



**University of
Sheffield**

**QUESTIONING PRODUCERS AND ACCEPTING CONSUMERS:
RESPONSES TO ARTIFICIAL INTELLIGENCE AMONGST LEGAL PROFESSIONALS**

Edward Alastair Walker

A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

The University of Sheffield

Faculty of Social Sciences

School of Management

October 2024

DECLARATION

I, the author, confirm that the Thesis is my own work. I am aware of the University's Guidance on the Use of Unfair Means (www.sheffield.ac.uk/ssid/unfair-means). This work has not been previously presented for an award at this, or any other, university.

STATEMENT OF COPYRIGHT

© Edward Alastair Walker, 2024, Questioning producers and accepting consumers: Responses to artificial intelligence amongst legal professionals, Sheffield University Management School, PhD Thesis ORCID 0000-0003-1399-6612

The copyright of this thesis rest with the author.

This copy of the thesis has been supplied on the condition that anyone who consults it is understood to recognise that its copyright rest with the author and that no quotation from the thesis and no information derived from it may be published without the appropriate permission or acknowledgement.

THE UNIVERSITY OF SHEFFIELD

ABSTRACT

**QUESTIONING PRODUCERS AND ACCEPTING CONSUMERS:
RESPONSES TO ARTIFICIAL INTELLIGENCE AMONGST LEGAL PROFESSIONALS**

Edward Alastair Walker

Doctor of Philosophy

October 2024

Recent advances in the field of artificial intelligence (AI) are having a profound effect on the contemporary workplace; reshaping the working practices of individual professionals and the structure of the firms they work within. Legal services are increasingly delivered through socio-technical systems underpinned by AI technology. This means many legal professionals now interact with AI on a daily basis in the delivery of their work. The ways in which experts respond to novel technologies can be hard to predict. While some show a passive acceptance of the technology and are happy for it to remain a 'black box', others question the technology to develop a deep understanding of how it works.

This research was motivated by a desire to better understand the experiences of legal professionals who use AI to deliver legal services to their clients. The research was phenomena-led and conducted using a critical realist case study design, with data generated at two UK commercial law firms, from a sample of 21 legal professionals.

The research contributes to our understanding of AI and the professions by exploring both the mental models through which legal professionals interpret AI, and the causal pathways that explain their use of AI. In doing so, the research reveals differences in the practices of 'producers' and 'consumers' of AI-enabled legal services, and the emergence of a novel 'liminal' role that has developed alongside them.

Keywords: Artificial Intelligence; Knowledge Work; Legal Services; Professional Service Firms

ACKNOWLEDGEMENTS

My PhD experience has required me to develop in several ways, many of which I could never have anticipated when starting my research. Such are the challenges of conducting research against the backdrop of a global pandemic. Reaching the end of this process was never a foregone conclusion, indeed it would not have been achievable without the support of several others.

I would like to thank Sheffield University Management School for their open-mindedness in accepting me as a mature student. The funding they extend to support PhD students, including myself, created the necessary time and space away from my professional commitments so I could focus on my academic studies. I also want to recognise the flexibility and support they have shown to students, during what has been a uniquely challenging period.

I would also like to acknowledge my supervisory team, Dr Chay Brooks, Professor Kamal Birdi and Dr Sarah Brooks. Between them they have provided the mixture of encouragement and candour I needed, while also retaining their good humour. A not insignificant feat as I recognise I must have been a frustrating student to supervise.

While the process of data generation was far from smooth, I also want to acknowledge the two law firms who granted me access to their professionals, during a period in which several other firms indicated their belief that conducting such research was impossible.

Finally, I would like to acknowledge the support of my partner Katherine, whose unstinting loyalty, helped me through several moments of self-doubt.

LIST OF CONTENTS

DECLARATION AND STATEMENT OF COPYRIGHT.....	1
ABSTRACT.....	2
ACKNOWLEDGEMENTS	3
LIST OF CONTENTS	4
1. INTRODUCTION	10
1.1 A typical day in the age of artificial intelligence.....	11
1.2 A short history of artificial intelligence	13
1.3 Public perceptions of artificial intelligence	16
1.4 Artificial intelligence and the UK economy	18
1.5 Defining artificial intelligence.....	21
1.6 AI-enabled professional and business services	26
1.7 Focal phenomenon: The use of AI-enabled legal services by legal professionals	31
1.8 My interest in AI-enabled legal services.....	33
1.9 Summary of research design	34
1.10 Summary of contribution	35
1.11 Thesis structure.....	36
2. LITERATURE REVIEW	37
2.1 My approach to conducting a literature review.....	37

2.2 Artificial intelligence in the professional services field	39
2.3 Review of the Information Systems literature	49
2.4 Trust and acceptance of artificial intelligence.....	58
2.5 Methodological implications of the literature review.....	70
2.6 The opportunity to make a contribution.....	74
2.7 Revisiting the literature during my research.....	76
2.8 Conclusion to the literature review.....	91
3. RESEARCH APPROACH AND STRATEGY	93
3.1 An introduction to phenomenon-based research.....	93
3.2 A phenomenon-based research strategy for AI-enabled legal services.....	94
3.3 Theoretical stance	99
3.4 Methodological principles and research design.....	108
4. CASE STUDY DESIGN AND METHODOLOGY	114
4.1 Assumptions of the case study design	114
4.2 Case study design	117
4.3 Wider methodological considerations	127
4.4 Methods for generation of data.....	139
5. ANALYTICAL STRATEGY	145
5.1 General principles for rigorous critical realist case study analysis.....	145

5.2	Theorising from case study data	147
5.3	Overview of the data analysis process	150
5.4	Process of thematic analysis	151
5.5	Process for identifying causality in case studies.....	169
6.	DATA ANALYSIS – THEMATIC ANALYSIS	174
6.1	Supplementary data	175
6.2	Dispositional theme development	188
6.3	Development of technological frame for AI-enabled legal services	231
7.	DATA ANALYSIS – PROCESS TRACING.....	237
7.1	Typical case selection	237
7.2	Within-case analysis of a typical producer case	243
7.3	Within-case analysis of a typical consumer case	253
8.	DISCUSSION AND RECOMMENDATIONS	264
8.1	Congruences and incongruences in the understanding of AI-enabled legal services	264
8.2	Extending the concept of consumers and producers of AI-enabled legal services	270
8.3	The use of AI-enabled legal services can be explained through two distinct causal pathways	274
8.4	Implications for organisational practitioners	278
8.5	Limitations of the research	281
8.6	Directions for future research.....	283

REFERENCES	286
------------------	-----

LIST OF APPENDICES

Appendix 1: Interview guide	306
Appendix 2: Information sheet	311
Appendix 3: Consent form	315
Appendix 4: Question list for organisational sponsors	316
Appendix 5: Participant instructions for reflective exercises.....	317
Appendix 6: Experiential theme summary.....	326

LIST OF TABLES

Table 1: Definitions of artificial intelligence.....	22
Table 2: The potential impact of artificial intelligence on a law firm.....	26
Table 3: Common categories of AI-enabled legal service	29
Table 4: Spectrum of human-AI interdependence.....	31
Table 5: How Cognitive Computer Systems challenge IS assumptions	52
Table 6: Heuristic devices chosen to generate data	74
Table 7: Dimensions of technological frames	79
Table 8: Steps in a ‘technology pipeline’	89
Table 9: Five phenomenon-based research strategies	97

Table 10: A phenomenon-based research strategy for AI-enabled legal services	98
Table 11: Different theoretical standpoints in case-based research	111
Table 12: Methodological implications of causal case study research	116
Table 13: Sampling criteria for the study of AI-enabled legal services	119
Table 14: Markers of rigour in critical realist case studies.....	146
Table 15: Four methods of theorising from case studies.....	148
Table 16: Critical realist approach to thematic analysis.....	154
Table 17: Theory building process tracing	169
Table 18: Contextual data relating to each participant.....	178
Table 19: Narrative timeline data relating to each participant	181
Table 20: Tetrad data relating to each participant	184
Table 21: Evaluation of AI-enabled legal services at National	187
Table 22: Evaluation of AI-enabled legal services at Global.....	188
Table 23: Shared structure and content of the technological frames of legal professionals relating to AI-enabled legal services	233
Table 24: Contrasting structure and content of the technological frames of producers and consumers relating to AI-enables legal services	235
Table 25: Summary of causal factors relating to each case	239
Table 26: Summary of Group A causal factors relating to each case	240
Table 27: Summary of Group B causal factors relating to each case	241

LIST OF FIGURES

Figure 1: Process steps in AI use and step-specific requirements in human resources	43
Figure 2: The process of knowledge work.....	48
Figure 3: Theory of Reasoned Action	59
Figure 4: Theory of Planned Behaviour	60
Figure 5: Technology Acceptance Model	62
Figure 6: Simplified Trust TAM model.....	66
Figure 7: Factors of trust development in human-technology interactions	67
Figure 8: A causal explanation of AI adoption using critical realist concepts	104
Figure 9: An example of forming explanatory statements from the thematic analysis	168
Figure 10: Theory building process tracing	172
Figure 11: Explanatory statement of dispositional theme one	190
Figure 12: Explanatory statement of dispositional theme two	205
Figure 13: Explanatory statement of dispositional theme three	217
Figure 14: Proposed causal mechanism explaining how producers validate the use of firm-developed AI software	249
Figure 15: Proposed causal mechanism explaining how producers validate the use of the external-developed AI software.....	249
Figure 16: Proposed causal mechanism explaining how consumers decide to use AI-enabled legal services as part of their professional practice	259

1. INTRODUCTION

*Lawyers with their suits and ties,
Now have AI by their side,
To analyse and sift through files,
And help them argue with their wiles.*

*Gone are the days of endless stacks,
Of papers, briefs, and legal acts,
Now AI can sort and search,
And find the needle in the perch.*

*Researching cases, parsing laws,
AI can do it without pause,
Speed and accuracy are key,
For lawyers in their legal spree.*

*AI can find the precedent,
And help the lawyer to present,
A winning argument in court,
With logic that can't be fought.*

*But still, the human touch is needed,
To navigate the law and lead it,
For lawyers bring their experience,
And judgment that's hard to sense.*

*So let the AI assist,
And make the lawyer's job less missed,
But in the end, it's the lawyer's skill,
That wins the case and pays the bill.*

A poem composed by ChatGPT (ChatGPT, no date) in response to my query, 'Write a poem about lawyers using artificial intelligence.'

In this introductory chapter I aim to provide a succinct overview of what this thesis is about and who is likely to benefit from the findings of the research. I start by providing a brief historical overview of artificial intelligence's (AI's) development from both an academic and commercial perspective. I then summarise the wider impact that technological change is having on the UK economy, highlighting the role of both AI and computerisation more generally. I then discuss the difficulty of defining AI and demarcating its boundaries, contrasting it with the related concept of automation, with which it is sometimes confused.

I then narrow my focus to the phenomenon of interest in this research, AI-enabled legal services, in the specific context of UK commercial law firms. I set out the academic reasons for choosing to study this phenomenon, highlighting the opportunity to make a theoretical contribution to the use of technology by knowledge workers - through identifying how AI is understood by them, and explaining the causal pathway that leads to AI use. In addition to highlighting the research questions I adopted in the research, I also explain my personal motivation for conducting the research and my position as a researcher. Based on this I then introduce the overall research design – a critical realist case study. I conclude the chapter by summarising the theoretical and practical contribution I believe my research has made.

1.1 A typical day in the age of artificial intelligence

“You wake up, refreshed, as your phone alarm goes off at 7:06am, having analysed your previous night’s sleep to work out the best point to interrupt your sleep cycle. You ask your voice assistant for an overview of the news, and it reads out a curated selection based on your interests. Your local MP is defending herself—a video has emerged which seems to show her privately attacking her party leader. The MP claims her face has been copied into the footage, and experts argue over the authenticity of the footage. As you leave, your daughter is practising for an upcoming exam with the help of an AI education app on her smartphone, which provides her with personalised content based on her strengths and weaknesses in previous lesson.

On your way to work, your car dashboard displays the latest traffic information, and estimates the length of your journey to the office, based on current traffic conditions and

data from previous journeys. On arrival, you check your emails, which have been automatically sifted into relevant categories for you. A colleague has sent you several dense legal documents, and software automatically highlights and summarises the points most relevant to a meeting you have later. You read another email, sent by your partner, asking if he can borrow your bank login details to quickly check something. On closer inspection you decide it is probably a fake, but still, you hesitate before deleting it, wondering briefly how the spammers captured his writing style so unerringly.

You have other things to worry about though, as you head to a hospital appointment. However, after a chest x-ray, you are surprised when the doctor sits you down immediately afterwards, explaining that you look to have a mild lung infection— you had expected it to take weeks before the results came back.

Your relief is short lived—a notification on your phone warns you of suspicious activity detected on your bank account, which has been automatically stopped as a result. You call the bank, and someone called Sarah picks up, and helps you order a replacement card. Except, you soon realise, Sarah is not human at all, just a piece of software which sounds just like a real person. You are a little unnerved you did not realise more quickly, but still, it got the job done, so you do not particularly mind.

After a quick detour to the local supermarket, where the products on the shelves have all been selected automatically based on previous customer demand, current shopping trends and the likely weather that day, you drive home. On your way back, your car detects signs that you are feeling slightly agitated, and chooses some music you have previously found relaxing. After dinner, you and your partner watch a film suggested by your TV, which somehow strikes just the right note for both of your normally divergent tastes. After dozing off, your house, predicting you are asleep by now, turns off the bathroom light and turns on the washing machine, ready for another day.” (UK Parliament, 2018)

The above vignette is taken from the executive summary of *AI in the UK: ready willing and able?* published by the UK Parliament in 2018. Its various examples of applied AI powerfully demonstrate the technology’s ability to infiltrate modern life, often without us being fully aware of it. A decade ago, most of the highlighted use cases would have been speculative,

but by 2018 the chosen examples reflected commercial technologies widely available to consumers living in the UK. The period since the report was published has seen AI make further inroads into our personal and professional lives.

While AI advocates argue that the technology will free individuals from the drudgery of routine work, thus creating time to focus on more novel and fulfilling tasks, there is growing evidence that AI is also starting to directly impact more complex tasks that were previously regarded as the sole preserve of humans. A recent example of this is ChatGPT (ChatGPT, no date), a highly versatile AI chat bot launched on 30 November 2022. Capable of, composing music, debugging computer code, generating images, and writing poetry (all in a matter of seconds); all the user has to do is type in everyday language what they want ChatGPT to do, as they would when writing an instruction for a human to follow. Within two months of its launch, it was reported that ChatGPT had been used by over 100 million people, making it the fastest growing consumer internet application to date (Milmo, 2023). This has led many researchers to suggest that AI is to be understood as producing a technological paradigm-shift effecting widespread change in organisational structures and the means of communication across multiple contexts. This means AI may ultimately be regarded as a general-purpose technology, comparable in its impact to the introduction of electricity and the internet; as opposed to merely a disruptive technology the impact of which, while significant, is typically experienced within a more restricted domain (Ågerfalk *et al.*, 2022). It is unusual for a single technology to have such a wide-ranging impact on our daily lives. This makes AI not just a technological phenomenon of interest to information systems researchers, but also a cultural and socio-political phenomenon of central importance to the fields of management and the broader social sciences (Lindgren and Holmström, 2020; Raisch and Krakowski, 2021). This highlights the importance of understanding AI from both an academic and practitioner perspective.

1.2 A short history of artificial intelligence

We are currently experiencing a 'boom' in AI, with high levels of interest in AI technologies amongst academics, business leaders, governments, the mass media, and consumers. However, it is important to note that the history of AI has been characterised by periods of

intense optimism, which were quickly followed by periods of disillusionment, often termed 'AI winters', during which interest in the technology waned significantly (UK AI Council, 2021)

The 1956 Dartmouth Conference is regarded as the symbolic birthplace of Artificial Intelligence, but earlier publications by McCulloch and Pitts (1943) on biological neural networks and Turing (1950) on intelligent machines, also made important contributions prior to this. Several projects emerged from the Dartmouth Conference with the shared goal of producing a computer program that could imitate human behaviours and thinking processes. This research attracted significant interest from both the US government and commercial organisations, leading to the period between 1956-1974 being regarded as the 'first golden age of AI' (Lee, 2020), with several management researchers of the time predicting that AI would become essential to modern management practices (Raisch and Krakowski, 2021).

Relatively quickly, two distinct approaches to developing AI emerged. The first adopted a mathematical approach to decision-making, based upon deductive reasoning and statistics. Rather than trying to imitate the way in which the human brain makes decisions, 'expert systems' were created that solved problems using programmed rules and knowledge provided by experts in a specific field. Hence the decisions expert system make can be seen to reflect what it has been previously taught. In contrast, the second approach sought to create AI that would reason in a way that was similar to the human brain. Drawing on knowledge from biology and psychology, AI systems were designed to mimic the behaviour of brain neurons. Rather than following predetermined rules, these 'neural networks' sought to process information based on their own previous learning, rather than basing it on expert human knowledge.

The initial progress made in AI research raised stakeholder expectations to unrealistic levels, leading to claims amongst some AI experts that machine intelligence would be developed 'within a generation'; the reality was very different (UK Parliament, 2018). Several high-profile project failures and a series of critical government reports led to research funds in the UK and US being withdrawn, and AI research in many institutions stalled. The period between the years 1974-80 was subsequently regarded as the first 'AI winter'. During this

period there was an effective moratorium on discussion of AI amongst management researchers. Instead, the consensus opinion was that while machines might be capable of performing routine operational tasks, complex managerial tasks that required intelligence were clearly the preserve of humans (Raisch and Krakowski, 2021). Humans, rather than machines, therefore, became the primary focus of organization and management scholars.

Despite these setbacks, AI and robotics research continued to thrive in Japan, which sowed the seeds for the 'second golden age of AI' (1980-87). During this period, significant progress was made in research areas such as neural networks; expert systems also became financially viable and were deployed in businesses to significant commercial effect (Lee, 2020). However, a second AI winter soon followed (1987-1993) as it became apparent that the early expert systems while useful, were only capable of solving niche problems; they were also severely hindered by insufficient computer processing power and data storage capabilities. Alongside these challenges facing AI, more general computing technologies were being developed that did not share the same limitations as AI-based systems. These technologies attracted the attention of researchers and organisations away from AI, and so progress in the AI field slowed once again.

The growth of the internet and mobile telecommunications during the 1990s proved an important stimulus for the re-emergence of AI, and specifically intelligent agents (computer programs that can perceive their environment and act in autonomous ways in order to achieve their predefined goals). Further developments in neural networks, which widened AI's range of applications, were underpinned by other significant technical innovations in both systems and components, including processors, sensors, algorithms, and software (Berente *et al.*, 2019). This meant that for the first time computers were unambiguously better than humans in computation and data processing tasks (Clifton, Glasmeier and Gray, 2020).

Unlike previous waves of AI research, many of these developments were funded by the private sector, rather than national governments. This led to several high-profile examples of commercially backed AI programs solving problems that were until then considered the preserve of humans. In 1997 Deep Blue (developed by IBM) defeated the reigning world chess champion; in 2011 Watson (IBM) won the US quiz show Jeopardy; in 2016 AlphaGo

(Google DeepMind) defeated the reigning world Go champion; and in 2022 there were claims that LaMDA (Google) had successfully passed the Turing Test - although this result is contested. The period from 1994 onwards is, therefore, considered the 'third golden age of AI'.

During this period management researchers have also re-engaged with the topic of AI leading to the publication of several highly influential books. Raisch and Krakowski (2021) drawing inspiration from Brynjolfsson and McAfee (2014), Daugherty and Wilson (2018) and Davenport and Kirby (2016) have, therefore, argued that AI should be repositioned at the centre of discussions about effective management. To date, however, a focus on the AI technology itself is essentially absent from management research. There is, therefore, a need for researchers to make AI visible and thus available for critical reflection. This is a necessary first step if AI is to be integrated into management theory, and for its role as an organisational agent (rather than a passive artifact) to be recognised (Raisch and Krakowski, 2021). My research, which focuses on the everyday interactions between AI and humans aims to contribute towards this goal, and through this help challenge some of the prevailing narratives surrounding workplace AI (Thompson and Graham, 2021).

1.3 Public perceptions of artificial intelligence

Awareness of AI amongst the general public has developed in parallel to that of government and business. Representation of AI in both the media and wider popular culture has had the effect of raising awareness of the technology, while at the same time presenting a frequently inaccurate picture of its capabilities, and the implications it may have (UK Parliament, 2018). This has led to public perceptions of AI being shaped by a narrative ecosystem comprised of both fictional and non-fictional accounts of AI.

The fictional accounts of AI typically display a common set of characteristics: a focus on embodiment in humanoid form; a tendency towards utopian or dystopian extremes; and a lack of diversity (Cave *et al.*, 2018). Non-fictional accounts are most often centred around the future of work debate – whether new technologies, such as AI, will create or destroy jobs, and if so, what sort of jobs. However, such reporting is often unbalanced, with

sensationalist media headlines seeking to grab their readers' attention and generate advertising revenue. The academic world has also proven itself prone to exaggeration, with researchers overstating the likely impact of AI, to attract research funding (UK Parliament, 2018). The result of this is that, in lieu of significant first-hand experience of AI, public attitudes towards AI are being shaped by perceptions of who controls the technology, what its purpose is, and who is likely to benefit (Cave *et al.*, 2018).

The recent growth in applied AI means the UK population is increasingly exposed to AI-enabled products and services on a regular basis. Through these experiences individuals are likely to become more aware of the material dimensions of different AI technologies, which will impact how the technology is interpreted. Such experiences may, therefore, be expected to play an increasingly important role in the shaping of individual attitudes, alongside the previously mentioned narrative accounts. Ultimately, individuals can be expected to create their own simplified mental model of AI through combining their material and symbolic knowledge of the technology. These models, which are sometimes termed *technological frames*, are important as they have been shown to influence individual behaviour towards technology (Orlikowski and Gash, 1994). Research has already indicated that diverse groups interpret AI in different ways to one another (Wang and Liang, 2024), hence, it is important that empirical research is conducted to generate knowledge of the structure and content of the frames held amongst specific populations, if we are to explain variations in the development and use of new technologies, such as AI.

However, the highly technical nature of AI means it may be unrealistic to expect people to understand how it works. This is especially true of AI applications that use algorithmic decision-making, a process whereby problems are solved through the use of a 'recipe' of predefined rules and steps, as an alternative to human judgement. The complex nature of these algorithms, sometimes even to expert audiences, leads to 'black box problems', where the embedded assumptions that underpin the algorithm lack technical transparency, making it impossible to explain how the outputs of AI are generated. The negative implications arising from this have already been seen in real-world settings, for example, Amazon's decision to scrap the use of an AI program used in employment selection, after it became apparent that it was displaying negative bias towards female applicants (Dastin,

2018). Such high-profile incidents when combined with a general lack of understanding about how the technology works, can lead to a lack of trust in the technology, which may act as a barrier to the adoption of AI.

Different solutions to these challenges have been proposed. Some proponents advocate the concepts of explainable AI (XAI), which requires that the output generated by AI must be understandable to human experts. An alternative view is that public engagement can build trust by focusing on the consequences of AI, rather than its inner workings; highlighting the benefits offered to humans, while also acknowledging the risks involved (UK Parliament, 2018). Understanding how trust (and mistrust) in AI develops is, therefore, an important topic for both researchers and practitioners interested in promoting the responsible use of AI.

1.4 Artificial intelligence and the UK economy

The use of AI is increasingly widespread across the UK economy and is seen by many policy makers and organisations as a key enabler of innovation, productivity, and economic growth (Brynjolfsson and McAfee, 2014; Susskind and Susskind, 2015). At the macro-level this is reflected in AI becoming a strategic pillar of government policy (HM Government, 2017); at the micro-level it is anticipated that AI will profoundly affect how work is done and by whom (Frey and Osborne, 2017).

Research cited in, *Growing the Artificial Intelligence Industry in the UK* (Hall and Pesenti, 2017), estimated that AI could add £630 billion to the value of the UK economy by 2035, and increase productivity in some industries by as much as 30%. The report highlighted commercial applications of AI having a significant impact across a range of industrial sectors including: financial services, customer services, manufacturing, and transport. The report also noted the growing importance of AI in traditional professions, such as accountancy, architecture, law and medicine. This analysis supports the widespread view amongst policy makers that AI's ability to use data in more efficient and innovative ways, has the potential to enhance most digital operations, products, and services in the UK economy .

There is also evidence that optimistic views about AI are held by a growing number of business leaders. In 2022 it was estimated that 15% of all UK businesses (432,000 companies) had adopted at least one AI technology, with a further 12% either piloting AI or indicating they were actively planning to adopt AI. While the rate of adoption varied across sectors, those displaying the highest rates of adoption - IT & Telecommunications (29.5%) and Legal (29.2%) – demonstrate that AI now underpins many commercial activities (Evans and Heimann, 2022).

These developments have been supported by the UK's digital tech sector, which was estimated to be worth £170 billion in 2015 and employed over 1.5m people. Given, the rate of growth in digital jobs from 2011-2015 was more than double the rate recorded for non-digital jobs (Hall and Pesenti, 2017) it is fair to assume that the digital sector will continue to grow in size and importance, relative to the wider economy. However, because AI intersects with several complementary general-purpose technologies, it is difficult to quantify AI's precise impact on the UK economy (Hall and Pesenti, 2017; Evans and Heimann, 2022); especially when consumers and businesses may be using AI without realising it.

The impact that AI (and technological change more generally) will have on employment and levels of inequality forms the core of the future of work debate (Brynjolfsson and McAfee, 2014; Faraj, Pachidi and Sayegh, 2018). The debate has seen several bold and competing claims made, often with limited empirical evidence available to substantiate them (Clifton, Glasmeier and Gray, 2020). Instead, predictions about the future of work tend to reflect expert opinion, for example, a survey of machine learning experts indicated a 50% likelihood of AI outperforming humans in all tasks within 45 years, and the automation of all human jobs within 120 years (Grace *et al.*, 2018). These types of predictions are often seized upon in media reporting, feeding into the wider public narratives of AI discussed previously.

The UK Parliament, (2018) highlighted the lack of consensus amongst experts when they tried to calculate the proportion of jobs at risk of unemployment because of AI. Based on research available at the time, the UK Parliament indicated a range of between 10-50% of jobs being at risk in developed economies, over the next 10-20 years. However, the same report also noted that historic concerns about mass unemployment resulting from earlier

waves of technological change, for example the introduction of robotic manufacturing, have failed to materialise. In reality a more complex picture arose, as while some jobs were lost, others were augmented, and new jobs emerged (Acemoglu and Restrepo, 2018).

The wide range of predictions also reflects differences in the methodologies used to calculate the figures. *Occupation-based approaches*, such as Frey and Osborne (2017), have estimated 47% of jobs in the US as being 'susceptible' to computerisation; the figure for the EU was calculated to be 54% (Bowles, 2014). In contrast, *task-based approaches* produce lower estimates, indicating that while certain tasks are at risk of computerisation, relatively few job roles are composed solely of such tasks; this means as few as 9% of jobs in developed economies may be considered at risk of computerisation (Arntz, Gregory and Zierahn, 2016).

It is also problematic that reports predicting future levels of employment lack precision with respect to how they define the causal technological phenomenon driving changing patterns of employment. For example, some reports reference 'the impact of 'AI'', others the impact of 'computerisation' (a much broader phenomenon, of which AI is just one element). It is, therefore, unsurprising that reports predicting the likely effect of computerisation on employment identify significantly larger effects than those restricted to the specific impact of AI. While a lack of definitional precision is understandable given that AI is typically embedded in other technologies (Evans and Heimann, 2022), we need to be careful not to overstate the likely impact of the technology.

There is also evidence that researchers and managers are using AI as a catch-all term for a range of technological changes taking place in the workplace (Clifton, Glasmeier and Gray, 2020). The magnitude of this problem is evidenced by the fact that in 2019, 40% of all European 'AI start-up' companies were found to use no recognisable AI technology in their products and services (Ram, 2019, cited in Dignam, 2020). This points to the importance of defining AI in a way that is flexible enough to recognise the different development focuses of AI (Lee, 2020), but which does not conflate it with separate phenomena, such as automation. For AI researchers interested in gaining an insight into public perceptions of AI, it also highlights the necessity of first determining what is understood to be AI, when respondents are invited to talk about their experience of the phenomenon.

1.5 Defining artificial intelligence

Given there is no widely accepted definition of the 'organic intelligence' demonstrated by humans and animals, it is perhaps to be expected that 'artificial intelligence' also remains a contested term (Ryan, 2020; Evans and Heimann, 2022). In their review of contemporary definitions of AI, Russell and Norvig (2016) highlight two important distinctions to consider when trying to define AI. The first distinction they make is between *behaviour* and *thinking*. Behavioural definitions tend to stress that AI refers to technologies that *act* like a human; whereas cognitive definitions state that what makes technology AI, is that it *thinks* like a human. Their second distinction relates to how the word *intelligence* is interpreted. Narrow definitions equate intelligence with *rational* thinking, meaning AI is seen to imitate human thinking if it is deemed to make logical decisions. Broad definitions of intelligence accept that not all human thinking can be considered rational, this would mean that to imitate human thinking, AI would need to demonstrate not just logical thinking, but also lateral thinking and intuitive thinking (Lee, 2020).

These differences become evident in the distinction between 'hard' and 'soft' AI. Hard AI, also called Artificial General Intelligence (AGI), focuses on the development of programs that can successfully perform any intellectual task that humans are capable of. Such programs would need to be capable of demonstrating intelligence in a range of different ways: problem solving, making judgements under uncertainty, planning, and communicating using natural language. While AGI is a developing area of AI research, this is not what currently available AI programs are designed to do. Soft AI refers to programs that are designed to mimic a limited range of human behaviour, typically thinking and problem solving. An example of this would be artificial neural networks, which use computer models to mimic the human brain. This narrower and more modest set of goals has helped soft AI research to make significant progress in well-defined areas, and has led to commercial applications of soft AI in contexts where decisions need to be made quickly using large, fuzzy data sets (Tredinnick, 2017). This explains why soft AI was the focus of this research, and, for the moment at least, means AGI is regarded by many as a distraction for researchers (Ågerfalk, *et al.*, 2022)

When discussing AI, it is also important to recognise that different categories of definition can be used, depending upon the aims of the researcher. Evans and Heimann (2022) distinguish three different ways of defining AI – *qualitative, usage-based and technology-based*. Qualitative definitions prioritise breadth, to help promote wider discussion of AI with non-expert audiences. This can be important when seeking to engage with users of AI who come from non-technical backgrounds. A limitation of qualitative definitions is that they lack precision, meaning they lack the specificity required for measurement. This means qualitative definitions are less useful when the goal of research is to classify or quantify AI.

Where quantification is the primary goal, a technology-based definition has the advantage of explicitly defining what technology is, and is not, classified as AI. This removes ambiguity and aids measurement. However, given the dynamic nature of AI development and AI's interconnected nature, it can be hard to accurately delineate technologies as AI or not AI; for example, when AI is embedded within a more general technology.

A usage-based definition focuses on *what* AI is doing, rather than *how* it is doing it. This approach offers insight into the different uses of AI, something which the other definition-types do not. However, by focusing on usage, the risk arises that those other technologies, which are being used to perform similar activities as AI, may mistakenly be labelled as AI.

The points raised above, are reflected in the different definitions of AI that have featured in recent reports on the subject. The selection of definitions provided in Table 1 highlight the widespread focus on soft AI, rather than AGI; while also demonstrating the use of different categories of definition, albeit with most definitions being qualitative in nature.

Table 1: Definitions of artificial intelligence

“AI is artificial mimicry of tasks and functions that would otherwise require human intelligence.” (Department for Business, Energy, and Industrial Strategy, 2017).

“Artificial Intelligence technologies aim to reproduce or surpass abilities (in computational systems) that would require 'intelligence' if humans were to perform them. These include: learning and adaptation; sensory understanding and interaction; reasoning and planning;

Table 1: Definitions of artificial intelligence (continued)

optimisation of procedures and parameters; autonomy; creativity; and extracting knowledge and predictions from large, diverse digital data.” (Engineering and Physical Sciences Research Council, 2022)

“Technologies with the ability to perform tasks that would otherwise require human intelligence, such as visual perception, speech recognition, and language translation...Our one addition to this definition is that AI systems today usually have the capacity to learn or adapt to new experiences or stimuli.” (UK Parliament, 2018)

This report uses “Artificial Intelligence” as an umbrella term to cover a set of complementary techniques that have developed from statistics, computer science and cognitive psychology. While recognising distinctions between specific technologies and terms...it is useful to see these technologies as a group, when considering how to support development and use of them. (Hall and Pesenti, 2017)

“While there are many permutations on the definition of AI, for the sake of brevity, we define it as a set of technologies that can imitate intelligent human behaviour” (KPMG cited in Clifton, Glasmeier and Gray, 2020)

“Artificial intelligence (AI) refers to machines performing cognitive functions that are usually associated with human minds, such as learning, interacting, and problem solving.” (Nilsson, 1971)

In my research I chose to adopt a qualitative definition of AI, to help facilitate my discussions of AI-enabled legal services with legal professionals. Conversely, because the classification and quantification of AI used by legal professionals was not a primary objective of the research, the advantage of adopting a usage-based or technology-based definition of AI was diminished.

Having decided to use a qualitative definition, the advantage of using a more colloquial definition, such as that provided by (Nilsson, 1971), became apparent. First, the definition used non-technical language that was capable of being understood by the participants in

this research; the definition also provided me with a straight-forward referent against which I could compare the participants' own definitions of AI. Second, the definition highlighted tasks AI is specifically associated with (learning, interacting and problem solving), which helped distinguish AI from other concepts the participants might be familiar with, for example, automation. Third, by not over-specifying the boundaries of AI, which would have restricted discussion of AI to specific use cases or specific technologies, the definition was both consistent with current AI terminology, and unlikely to become outdated during the research process. Fourth, by not focusing on specific technologies the definition provided a degree of flexibility when generating data, such that the legal professionals could focus on what AI meant to them in the context of their work, rather than trying to fit their experiences within a researcher-defined framework, which would have meant the data generation process was less participant-led.

Distinguishing artificial intelligence from automation

Alongside AI, automation is another concept of significant relevance to the future of work debate. The two concepts are often spoken about together and can at times get conflated with one another (Clifton, Glasmeier and Gray, 2020). In order to separate the two, we need to return to the previous discussion about what is *intelligence*.

Automation can be defined as technology that allows a process to be performed without (or with minimal) human assistance (Groover, 2020). The aim of automation is to replicate tasks previously performed by humans, in ways that are more efficient, predictable, and reliable. Automation becomes feasible when a task (or role) no longer requires human intelligence to be performed to the required standard. Automated tasks are characterised by deterministic ('if-then') processes, this means a given input, always produces the same output, using the same process e.g. a calculator determines five multiplied by ten always equals fifty.

For automation to be possible, the sequence of actions required to generate the desired outcome needs to be precisely defined; this typically requires close collaboration between humans and machines. This can be seen where over time humans as the users of machines, are able to identify rules or develop models that robustly solve problems or perform tasks,

to a level of accuracy that meets (or exceeds) the required standard. By integrating the rules or model into the process it then becomes possible to remove the need for human intervention on an ongoing basis, and the task becomes automated with the machine performing it independently (Raisch and Krakowski, 2021).

The ability to perform a predefined task to a high standard does not, however, make a machine intelligent. All the machine is doing is replicating a process which was developed through human intelligence; the machine could not have developed the process itself in isolation. This highlights significant limitations of automated systems: they do not possess the intelligence required to learn from previous experience; move away from a predefined process or; change the outcome of a process, even when this is desirable.

In contrast to automation, AI systems possess the ability to interpret data, learn from the data, and use what is learnt to adapt their own processes, to achieve specific goals. AI's ability to perform such activities is typically underpinned by machine learning techniques that, as a result of the continued increase in the outer limits of data processing and data storage, are now capable of processing large amounts of unstructured data, using algorithmic models which reflect the AI's previous experience, rather than the explicit programming of a human. The data can then be used to learn, predict, and perform tasks. This means to outside observers an AI system can appear like an automated system, as both are performing tasks without human intervention. However, under closer examination the two are quite different, as while the automated system is fixed and will not change without human intervention; the AI system possesses the 'intelligence' to monitor and refine what it is doing in a dynamic way, albeit the changes it makes may not be readily visible to onlookers.

To summarise, the rules and models that in a traditional automated process would be developed by a human, based on their knowledge and experience, are now generated using AI to a level of sophistication that is often beyond human understanding. This means the processes developed by AI can now outperform those developed by a human, and will continue to improve over time, provided the data necessary for learning to take place is made accessible.

1.6 AI-enabled professional and business services

The professional and business services (PBS) sector is an important constituent of developed economies, such as the UK. The sector includes a diverse range of knowledge-intensive industries that produce specialist support and advice to private and public organisations throughout the economy. In the UK, PBS accounted for £186 billion gross value added and in 2016 employed 4.6 million people, equivalent to 13% of the UK workforce (HM Government, 2017). Most people working within the sector can be described as knowledge workers – people who ‘think’ for a living, applying their theoretical and analytical knowledge to generate knowledge that can be sold as a product or service (Bechky, 2006). Internationally, the sector is also an important exporter, accounting for 27% of the UK’s services exports (valued at £66 billion).

While various organisational forms exist within the PBS sector, the most common form is the professional service firm. Empson *et al.* (2015) identify four characteristics that must be possessed, for an organisation to be considered a professional service firm (PSF). While my research focuses on the professionals working within PSFs, it is still important to understand the characteristics of PSFs, as these shape the context within which these professionals are working. The four characteristics provide a framework to think about how AI-enabled services might impact the wider sector and the individual professionals working within it. Table 2 highlights some of the ways in which AI might impact a law firm, through reference to the four defining characteristics of a PSF.

Table 2: The potential impact of artificial intelligence on a law firm

Primary activity: Applying specialist knowledge to create customized client solutions.

PSFs are knowledge-intensive rather than capital-intensive organisations. The bespoke nature of their work requires an intensive relationship between professionals and their clients. For example, lawyers typically produce several drafts, responding to client feedback, to ensure the resulting contract reflects the client’s idiosyncrasies. AI could impact this primary activity by augmenting the work of lawyers by using historic client data to identify insights that legal professionals might not identify themselves, thus allowing them to create novel solutions for clients.

Table 2: The potential impact of artificial intelligence on a law firm (continued)

Knowledge: Professional knowledge and knowledge of clients is a core asset.

Professionally accredited knowledge is only one aspect of PSF knowledge. Methods for acquiring and sharing knowledge, and knowledge held about clients are also assets, as these enable firms to apply their specialist expertise. For example, while a lawyer’s legal knowledge is fundamental to performing their role, it is only through combination with knowledge of their client that they can offer bespoke advice. AI’s ability to process and analyse vast amounts of data means law firms that are able to capture tacit knowledge from their lawyers, can use AI to aggregate this knowledge and create value from it in ways that were not previously possible.

Governance: Extensive individual autonomy combined with contingent managerial authority; producers own or control core assets.

Qualified professionals are granted significant autonomy by their professional regulators and employers, so they can act in the best interests of their clients. This means managerial control tends to be lower in PSFs than in other sectors. Legal processes need to be formalised if they are to leverage AI technology in a consistent and efficient manner. This poses a challenge to the high-level of autonomy that individual lawyers have historically enjoyed; strengthening the power of those who own and are responsible for the management of AI within the firm.

Identity: Core producers recognise one another as professionals and are recognised as such by clients and competitors.

Loose governance means PSFs rely on a shared sense of professional identity to guide individual behaviour. Traditionally, professional identity develops through formal training and the award of qualifications, but increasingly it also reflects the professional norms of the PSF. The introduction of AI-enabled legal services has the potential to disturb the content of legal education and the norms of professional practice, thus causing individual lawyers to re-evaluate their sense of self.

Alongside the growth of the PBS sector and the numbers of knowledge workers it employs, a ‘technization’ of professional practice has also taken place. This process involves the use of ‘epistemic technologies’ that professionals are increasingly required to use to construct

knowledge (Anthony, 2018). The use of such technology is reshaping long-established professional processes and their output, in terms of speed, accuracy and novelty. It is, therefore, widely recognised that AI provides unprecedented opportunities to innovate business services and the organizational models that deliver them; but AI also raises challenges for businesses (Berente *et al.*, 2019). These include the development of AI-enabled services, underpinned by processes that require humans and AI to interface with one another. Technological change has long been understood to have the potential to challenge the expertise and status of knowledge workers (for example, Nelson and Irwin, 2014). In response to such changes, knowledge workers have historically been shown to respond by seeking to develop a deep understanding of any technology that underpins the services they provide, in order to assure the quality of their work and retain their expert status (Bailey and Barley, 2011). However, more recent research suggests that this response is not always seen following the introduction of AI. Instead a mixed set of responses has been observed, with some workers continuing to question the technology before using it; while others accept its introduction despite not understanding how it functions (Anthony, 2021). How professionals develop trust in AI, and whether they are justified in doing so, is an important process to understand, given the trust wider society places in the judgement of professionals. It is, therefore, important that research that explains how knowledge workers respond to AI is undertaken, in order to help guide organisations through this complex environment, thereby helping to ensure the ongoing success of the PBS sector and protect the interests of those who rely on their services.

AI-enabled legal services

As part of the PBS sector, the UK legal services sector had estimated revenues of £37 billion in 2019. Historically the legal services sector was characterised by high human capital intensity and low specific capital intensity (Smets *et al.*, 2017). This led to increasing recognition of the potential for technology to transform both the structure of the sector (Hinings, Gegenhuber and Greenwood, 2018; Kronblad, 2020), and the delivery of legal services (Katz, 2013; McGinnis and Pearce, 2014). More recently both academic and commercial attention has focused on the potential impact that AI could have on the legal sector. The commercial importance of AI is reflected by the finding that in the UK, AI adoption levels in the legal industry are significantly higher than other sectors of the

economy (Evans and Heimann, 2022). However, different segments of the UK legal market display varying levels of technological maturity, with AI primarily being adopted by large commercial firms who provide legal services to corporate clients (The Law Society, 2019).

At a global level, the Legaltech sector has grown rapidly; in 2018 the value of Legal AI was estimated at \$3 billion, and was predicted to grow to over \$37 billion by 2026 (Zion Market Research, 2019). In addition to being relatively early adopters of AI, law firms are widely regarded as being archetypes of the professional service firm (von Nordenflycht, 2010). Thus, making them a suitable context within which to research the impact of AI on professionals.

Table 3: Common categories of AI-enabled legal service

Category	Task description	Benefit	Example AI Products
Due diligence	Identify pertinent information amongst large data sets for review by legal professionals. This includes reviewing contracts, conducting legal research and electronic discovery.	Speed and accuracy	eBrevia, iManage, Kira Systems
Prediction	Forecast the likelihood of different litigation outcomes.	Accuracy and Insight	Casetext, Intraspection, Premonition,
Legal analytics	Analyse historical data to provide insights that can assist lawyer decision-making	Speed and Insight	Lex Machina, Loom Analytics, Solomonic
Document creation	Create complete documents based on the client needs	Speed and Accuracy	Clio, Neota
Electronic billing	Record and calculate lawyer billable hours and generate invoices.	Speed and Accuracy	Brightflag, Smokeball

Source: Adapted from Faggella (2020)

AI has been used within the legal sector to fulfil several different purposes (see Table 3 above). The primary driver of AI adoption is efficiency, with AI-driven processes capable of performing certain tasks more quickly and accurately than legal professionals (The Law Society, 2019).

The processes which underpin AI-enabled legal services can be numerous and complex, and are dependent on the service being offered. Each process typically involves several discrete steps, with the degree of interdependence between AI and legal professionals varying according to the precise nature of the task. The term ‘legal professionals’, rather than ‘lawyers’ or ‘solicitors’ is used when discussing AI-enabled legal services, as these services invariably require input from a combination of practising lawyers and professionals from other specialisms. This means within law firms there is range of different individuals whose roles bring them into direct contact with AI-enabled legal services (Armour and Sako, 2020). More recently this has led to a distinction being made between consumers and producers of AI-enabled legal services (Armour, Parnham and Sako, 2022), with *consumers* referring to legal fee earners who use AI in their work, whereas *producers* are those individuals who are responsible for the technology that augments AI-enabled legal services. While consumer-producer distinction is conceptually useful, further empirical evidence is required in order to evaluate the extent to which these two categories accurately capture organisational reality; or whether further categories are required to explain the different ways in which legal professionals interact with AI-enabled legal services.

Indeed, it has been suggested that human-AI interdependence can be understood as a spectrum (Table 4) ranging from task substitution, through task augmentation and task assemblage, beyond which interdependence ceases and tasks remain the exclusive preserve of humans (Rai, Constantinides and Sarker, 2019). While AI was initially anticipated as a technological substitute that would drive the automation of processes, many of the most promising use cases of AI at the present moment, see people and computers working together as ‘superminds’ (Malone, Rus and Laubacher, 2020; Wilson and Daugherty, 2018).

The use of AI within law firms to date sits at various points along this spectrum, with task augmentation (from AI-to-human and human-to-AI) frequently taking place (Armour, Parnham and Sako, 2022). This shift towards AI-enabled legal services means an

understanding of human-AI interaction, and how human-AI relationships develop, is of importance to both the academic study of the professions, and practitioners working within the sector.

Table 4: Spectrum of human-AI interdependence

Task substitution	AI used to substitute for the role of human in performing a task. No human intervention is required.
Task augmentation	Humans and AI augment one another to perform a task. This typically involves the task being partitioned. This may take the form of the AI augmenting the human e.g. a legal professional uses AI to identify documents that meet certain criteria, prior to reviewing them in more detail. Alternatively, humans can augment AI e.g. the legal professional explains to their client how the AI arrived at a specific recommendation, relating to the strategy of a case.
Task assemblage	Humans and AI work dynamically together to perform an emergent task, without the task being partitioned. Examples of such assemblages are still rare but can be seen in domains such as medicine where surgeons use AI-powered robots to perform minimally invasive surgery.
Human tasks	Humans undertake the task without any involvement from AI, as they retain a comparative advantage in performing it.

Source: Adapted from Rai, Constantinides and Sarker (2019)

1.7 Focal phenomenon: The use of AI-enabled legal services by legal professionals

My research was undertaken to develop understanding about the use of AI-enabled services by legal professionals working in UK law firms. As an embryonic phenomenon limited conceptualising and theorising about AI-enabled legal services has taken place to date, and there remains a lack of empirical research of the phenomenon (Sako and Parnham, 2021). It has, however, been suggested that the material characteristics of AI challenge several assumptions relating to how users interact with information systems (Schuetz and Venkatesh, 2020; Anthony, Bechky and Fayard, 2023), which in turn underpin existing

theories of technology acceptance. This meant there was a need to undertake research to determine whether the use of AI-enabled legal services follows a similar process to other technological innovations, or represents something novel and theoretically distinct.

In addition, the distinct context of professional service firms means legal professionals work in an environment that is not reflective of the wider economy, suggesting the generalising of findings from IS and management research undertaken in other contexts is problematic (Brock, Leblebici and Muzio, 2014). The literature relating to expertise within organisations, which is very relevant to knowledge-intensive organisations, such as law firms, suggests that non-experts, have been found to trust technological tools without necessarily understanding how they work. In contrast, it has been theorised that knowledge workers, such as legal professionals, seek to understand the tools they use in order to preserve the perceived quality of their work and their expert status (Bailey and Barley, 2011).

Professional membership has also been shown to impact perceptions of what is considered legitimate behaviour, with individuals strongly attached to existing norms of conduct (Anteby, Chan and DiBenigno, 2016). Hence, in comparison to other workers, professionals can require stronger external pressures (Anthony, 2021) and peer guidance (Bourmault and Anteby, 2023) before they are willing to reassess their working practices.

To date there is limited empirical research to explain the ways in which legal professionals have adapted their practices in response to the introduction of AI, meaning it is unclear when and how such change takes place. My research offered the opportunity to investigate this phenomenon empirically, and generate a plausible explanation for AI use amongst legal professionals by focusing on identifying the causal factors capable of triggering the use of AI by legal professionals, and the causal mechanisms through which these factors operated.

Research questions

Taking inspiration from the 'researcher as detective' approach to studying the professions (Sherer, 2019) my research questions, emerged from my initial review of the extant literature (discussed in Chapter 2.6), before being refined in response to the data I was able to generate (discussed in Chapter 5.4). This led to the research focusing on two overarching research questions.

Research question 1: How are AI-enabled legal services understood by the legal professionals that use them?

This question sought to identify the assumptions and expectations that legal professionals had about the role of AI within AI-enabled legal services, and the factors that could explain how these views were formed. Taken together these findings aimed to provide insights into the structure and content of the mental models that legal professionals developed in relation to AI-enabled legal services.

Research question 2 'What explained the use of AI-enabled legal services amongst these legal professionals?'

Answering this question, required the identification of both the causal factors and mechanisms that together could explain the use of AI-enabled legal services by the participants in the research. It also provided the opportunity to identify variations in the ways in which different groups of legal professionals – specifically *producers* and *consumers* (Armour, Parnham and Sako, 2022) – became users of AI-enabled legal services.

1.8 My interest in AI-enabled legal services

Running in parallel to the growing academic interest in AI-enabled legal services, was my own professional interest in the phenomenon, which led me to undertake this research. The legal sector has been near ever-present throughout my life, despite the fact that I have never studied or practised Law. My father was a lawyer for over forty years, a period during which the profession changed significantly, in large part due to the introduction of new business models and technology. Prior to my PhD studies, I had worked within the legal sector for over a decade, first within the human resources function of a large commercial law firm and more recently as a consultant to several leading international law firms. This experience made me a witness to the introduction of AI-enabled legal services and exposed me to the general discourse that circulates around the phenomenon. The primary focus of my professional work has been helping individual legal professionals develop the skills and capabilities they need to drive a successful career. It was through this work that I recognised the potential of AI-enabled legal services to have both positive and negative

effects on those who interacted with them; and the impact this had on the behaviour of legal professionals. I was, therefore, motivated to learn more about the phenomenon, both academically and professional, in order to help legal professionals thrive in the era of AI-enabled legal services.

1.9 Summary of research design

Given the relatively embryonic nature of AI-enabled legal services, I adopted a phenomenon-based approach in my research; focusing on distinguishing and exploring the phenomenon, with a view to developing a theoretical explanation for any puzzles that emerged (von Krogh, Rossi-Lamastra and Haefliger, 2012). Critical realism (Bhaskar, 1997, 1998) provided a suitable theoretical framework for the research, with its focus on making sense of a phenomenon through identifying the social processes related to it, and explaining the causal mechanisms underpinning these processes. Critical realism's view of causality as a combination of individual agency embedded within a wider social structure was well-suited to researching a socio-technical phenomenon, where the decision to adopt AI-enabled legal services, would reflect the decision-making of individual legal professionals, the affordances of the technology, and structural factors within the wider context in which AI use took place.

A case study approach (underpinned by critical realist principles) was chosen as the research design. This reflected its suitability to investigate a contemporary phenomenon within its real-life context, and its ability to generate causal explanations that could be generalised within a small, bounded population. The research generated data for analysis through a combination of written exercises and semi-structured interviews with legal professionals, supplemented with my own first-hand observations of AI-enabled legal services. Analysis of the data was first undertaken using a critical realist thematic analysis (Wiltshire and Ronkainen, 2021) to make both the human and the technology within AI-enabled legal services visible (Adams and Thompson, 2016). Following this, the data was further analysed using causal case study methods (Beach and Pedersen, 2016), to identify the combination of causal factors and mechanisms responsible for the adoption of AI-enabled legal services.

1.10 Summary of contribution

The questions explored through this research, which focused on the use of AI-enabled services by legal professionals within the UK legal sector, generated findings that contribute to our theoretical and practical understanding of the phenomenon in the following ways:

1. Through investigating how legal professionals understood the role of AI in their professional practice, the research identified both congruences and incongruences in the structure and content of the technological frames (Orlikowski and Gash, 1994) that different groups of legal professionals developed through their experience of AI-enabled legal services. The findings suggest a common five-dimension structure to the technological frames of legal professionals. This structure differs from more generalised theoretical models (Spieth *et al.*, 2021), with a novel dimension called *expert influence* found to replace the theorised dimension of *supervisor influence*. This revised structure provides a parsimonious way to understand the different factors which influence the decision-making of legal professionals towards AI-enabled legal services.
2. Analysis of the different roles undertaken by legal professionals in this research led to the identification of a new category of legal professionals, not previously accounted for in the extant research. Hence, while the empirical findings support earlier conceptualisations of *producers* and *consumers* of AI-enabled legal services (Armour, Parnham and Sako, 2022), the findings also suggest that a new category called *liminals*, needs to be added to the existing model of the technology pipeline, if it is to accurately reflect current organisational practice within UK law firms.
3. Through tracing the causal factors and mechanisms through which legal professionals came to use AI, the research provides a localised theoretical explanation for the use of AI-enabled legal services by both consumers and producers. Through the application of process tracing, it is demonstrated that the use of AI-enabled legal services can be explained through two distinct causal pathways, with evidence of questioning practices in the producer pathway and accepting practices in the consumer pathway (Anthony, 2018). The pathways also highlight distinctive ways in which producers and consumers develop their trust in AI-enabled legal services; reflecting the different trusting bases

they find most salient. Together, these findings provide further empirical evidence that in certain contexts professionals will incorporate novel technologies into their professional practice, despite not understanding how the technology works (Anthony, 2021); while also highlighting that the precise process through which this is achieved, is highly localised and varies amongst legal professionals with different job roles.

1.11 Thesis Structure

In addition to this introduction the thesis is comprised of seven further chapters. Chapter 2, the *literature review*, starts by summarising both the theoretical frameworks and empirical research that informed the development of my initial research questions. The review then explains how the literature relating to the use of AI by professionals has continued to develop during my research; and the impact this has had.

Chapter 3, *research approach and strategy*, outlines the research strategy used to investigate AI-enabled legal services. It explains my reasons for adopting a critical realist case study design and the methodological principles that follow. A more detailed discussion of the methodological decisions relating to case study design, including case selection and methods for generating data follows in Chapter 4, *case study design and methodology*.

Chapter 5, *analytical strategy*, details the two approaches used to analyse the data generated - critical realist thematic analysis (Wiltshire and Ronkainen, 2021) and process tracing (Beach and Pedersen, 2016). Chapter 6 reports both the key findings of the thematic analysis and details the technological frame that was developed from this data. Chapter 7 reports the within-case analysis of typical *producer* and *consumer* cases, used to identify the causal pathway that led to AI use amongst different groups of legal professionals.

The final chapter, *discussion and recommendations*, presents the theoretical contribution the findings make to our understanding of AI-enabled legal services. Following this, the chapter highlights potential implications for organisational practitioners, the limitations of the research and further opportunities for research.

2. LITERATURE REVIEW

2.1 My approach to conducting a literature review

My review of previous research relating to AI-enabled legal services is written using a narrative approach, meaning it reflects my subjective understanding of this emerging phenomenon at the time of writing. Over the course of my research, interest in AI-related phenomena has increased significantly across different fields, with systematic literature reviews of artificial intelligence (AI) beginning to be published within the Information Services field (for example, Collins *et al.* (2021); Çelebi (2021)). Elsewhere the volume of AI-related research remains more limited, meaning it was premature for me to employ more systematic methodologies to review the AI-related literature within the field of professional services.

I adopted a hermeneutic approach to conducting my literature review, which is more iterative than other approaches. Rather than seeking to analyse and summarise the extant research in a methodical way, the approach places a greater emphasis on identifying the meanings and assumptions found within literature, and the contexts within which it exists. This meant prioritising the development of my understanding as a researcher, rather than seeking to produce a review that could be regarded as objective and replicable (Boell and Cecez-Kecmanovic, 2014). As the researcher, I was, therefore, situated within the review process, meaning the way I conducted the review was influenced by my beliefs and past experiences (Ramalho *et al.*, 2015).

This means the content of my review is idiosyncratic, and reflects my professional background as an organisational psychologist, who has worked extensively with individuals in professional service firms. These experiences led me to develop an interest in AI-enabled legal services; and having provided training and coaching to legal professionals for several years, informed my decision to focus my research on understanding how these services affect the working lives of individual professionals. I was, therefore, particularly interested in identifying research that considered the impact of the phenomenon at the individual

level. My professional background also exposed me to the use of biopsychosocial models to understand individual health and well-being (Karunamuni, Imayama and Goonetilleke, 2021). Having seen the value of applying interdisciplinary frameworks to understand complex phenomena and human behaviour, I was, therefore, predisposed to look for insights about AI-enabled legal services across different academic disciplines, rather than being wedded to a specific perspective. This is reflected in my initial literature review, which was completed relatively early in my studies, and which draws upon research from the professional services field and the domain of information systems (IS) research, which are themselves underpinned by theories from a range of different academic disciplines.

Ultimately, the aim of my initial review was to demonstrate the scarcity of research about AI-enabled legal services, and identify sensitising concepts (Blumer, 1954) that had the potential to guide me in my research. I was, therefore, not expecting to identify an existing theoretical framework that I could use in a prescriptive fashion to research AI-enabled legal services (Bowen, 2006). As I progressed my research and began to generate and analyse my data, I returned to the extant literature in anticipation of finding additional concepts and theories that would help me to better interpret and understand my data. I also took the opportunity to consider more recent literature pertaining to AI-enabled legal services that had not been published at the time I conducted my initial literature review. Where relevant, this allowed me to re-engage and deepen my understanding of the theories and concepts identified in my initial review of the literature.

Literature review structure

While the initial literature review is presented in a linear fashion to make the narrative easier to follow, my review method involved moving back and forth between different disciplines, as my knowledge of the phenomenon grew. The initial literature review begins with a discussion of AI-research within the professional services field (section 2.2).

Reviewing the different research streams through which the field has sought to understand AI, gave me the opportunity to evaluate the literature and identify prevailing discourses surrounding the topic. This provided important context and helped me to understand how my research interests related to the extant literature (Charmaz, 2006). The review also revealed several sensitising concepts – institutional logics, professional identity and models

of professionalism – that offered me different perspectives to think about AI-enabled legal services, and which had the potential to help me develop a theoretical explanation of the empirical data my research generated (Bowen, 2006).

Importantly, this initial review of the literature also acted as a stimulus to my own thinking, as it allowed me to identify under-researched aspects of AI within the professional services field – how AI is understood by professionals; and how professionals interact with one another and AI technology to deliver AI-enabled services – where there was the potential to make a theoretical contribution that would act as a justification for my research.

As research on AI has taken place across different disciplines, the next section (2.3) of the initial review details the concepts and theories I engaged with to refine my thinking about AI-enabled legal services. My starting point for this was research focused on *IT and individuals*, a key research stream within the domain of information systems (IS), which in addition to its own theoretical base, also draws heavily on theories from psychology and sociology (Lim *et al.*, 2013). This helped me identify further sensitising concepts, specifically theories relating to technology adoption and trust, that highlighted aspects of AI-enabled legal services that it would be valuable to explore (section 2.4).

Reviewing this literature also enabled me to identify the assumptions that underpin many of the core theories of IS research, which AI-enabled legal services had the potential to challenge (Schuetz and Venkatesh, 2020); and alerted me to important methodological issues that AI research should ideally address (Bailey and Barley, 2020). The impact this had on the formulation of my initial research questions, and my approach to generating data, is discussed towards the end of the initial review section (sections 2.5 and 2.6). My further engagement with the literature during the course of my studies is summarised in section 2.7, followed by a general conclusion to the literature review process (section 2.8).

2.2 Artificial Intelligence in the professional services field

On commencing my studies, my knowledge of AI reflected my professional experience, rather than the extant academic research. To develop my understanding of the professional services field I consulted an overview of the field conducted by Brock, Leblebici and Muzio

(2014), which categorised the major streams of research relating to professionals and the organisations they work in. Their analysis gave me a useful vantage point to identify the wider discussions that contemporary phenomena, such as AI-enabled legal services, had been integrated within. Noting that the impact of technology on professionals (and the professions) was not identified as a distinct conversation, I recognised that I would need to look for research relating to technological phenomena across different research streams, the most relevant of which I discuss below.

Institutional theory & logics

Research relating to *Organisational models and structures* focuses on understanding what makes professional service firms (PSFs) distinct from other organisations. Engaging in this literature allowed me to see how the field has tracked changes in the structure of professional organisations, in response to a number of exogenous forces such as globalisation (Faulconbridge and Muzio, 2017), deregulation (von Nordenflycht, 2014) and technological change (Hinings, Gegenhuber and Greenwood, 2018).

In seeking to explain the effects of these forces, researchers have drawn heavily on the organizational studies literature, including institutional theory (DiMaggio and Powell, 1983) and institutional logics (Friedland and Alford, 1991). Institutional theory predicts that organisations within a given industry sector, will show a high-level of homogeneity in both their structure and practices; this was something I had observed first-hand working professionally with UK and international law firms over several years. I, therefore, quickly recognised the potential value of the concept of coercive, normative and mimetic pressures in helping me understand how AI-enabled legal services had developed within the legal sector.

The concept of institutional logics, as a set of norms and beliefs that can regulate the behaviour of both individuals and organisations (but in a non-deterministic way) also appealed to me. My psychology background has meant I consider humans to have agency over their behaviour, while also recognising the role that social factors can play alongside this. Hence while I was receptive to the idea that behaviour could be guided by logics that reflected wider structural factors, it was important to me that the theory acknowledged

that logics can be multiple, can potentially conflict, and require individual interpretation; thereby allowing variations in behaviour to be seen within a shared context (Reay *et al.*, 2017). Therefore, when seeking to understand the behaviour of professionals using AI-enabled legal services, institutional logics reminded me to consider both the direct impact that technology can have on behaviour, and the structural context that the behaviour takes place within (Goto, 2021).

Institutional theory and logics have already been used to analyse how the digitalisation of the legal profession has impacted the ways in which legal services are delivered (Kronblad, 2020), and how AI has transformed the business models of both established law firms and newer market entrants (Brooks, Gherhes and Vorley, 2020; Callegari and Rai, 2022). While this research focused on the impact of technology at the organisational level (as opposed to the individual-level focus of my own research), it caused me to reflect on the importance of conducting my research using methods that would permit the generation of data relating to structural factors; and that in my analysis, institutional logics offered a useful interpretive lens to understand the different ways in which legal professionals might think about and interact with AI-enabled legal services.

This stream of research also alerted me to the challenge of defining artificial intelligence (discussed in Chapter 1). In using the term *digitalisation* Kronblad's (2020) research sidestepped this problem, and instead focused on the adoption and increased use of all computer technologies by PSFs. In acknowledging that digitalisation covers a range of technologies, the research is unable to offer specific insights into AI-enabled legal services, as it does not distinguish AI from the other technologies that law firms use. In contrast to this Brooks, Gherhes and Vorley (2020) discuss the difficulties in defining AI, before adopting a still broad perspective that considers the impact of Legaltech, which includes both automation and AI-based tools. While this research is more focused than Kronblad's (2020), AI and automation should arguably be treated as separate phenomena, each with their own distinguishing features (as discussed in Chapter 1). I, therefore, concluded that conducting my research with a narrow focus on AI-enabled legal services, would enable me to make the phenomenon more visible. Thereby allowing me to identify and investigate its distinct features, with the ultimate goal of better understanding how legal professionals

interact with AI-enabled legal services, in comparison to other technologies (Chatterjee *et al.*, 2021).

Managing professionals and professional identity in the context of artificial intelligence

Research on *Micro-organisational issues*, foregrounds factors that affect the management of professionals. This broad category of research includes disparate topics including: leadership practices, management systems, career paths, and identity work (Brock, Leblebici and Muzio, 2014). Historically this research stream received less attention than research focused on the context and tasks of professional service firms, despite many management issues being expressed in unique ways across different professions (Suddaby, Greenwood and Wilderom, 2008). However, the same exogenous forces that stimulated research within the *Organisational models and structures* stream, have inspired research at the micro-level, with the impact of technology on the work of individual professionals one facet of this (see for example, Nelson and Irwin, 2014). A more detailed discussion of the literature relating to how knowledge workers respond to technology can be found in section 2.7.

Theoretical discussions about the impact of AI on the practice of law initially took place within legal journals. This research generated important insights about the type of legal tasks AI has the capability to displace or augment, and what the limits of legal computation were likely to be (Katz, 2013; McGinnis and Pearce, 2014; Markou and Deakin, 2020). Within a few years the debate had evolved to discuss commercially available software tools and actual use cases (Alarie, Niblett and Yoon, 2018). Building on these insights Armour and Sako (2020) developed a process map, which outlined a workflow to solve a legal problem using AI (see Figure 1).

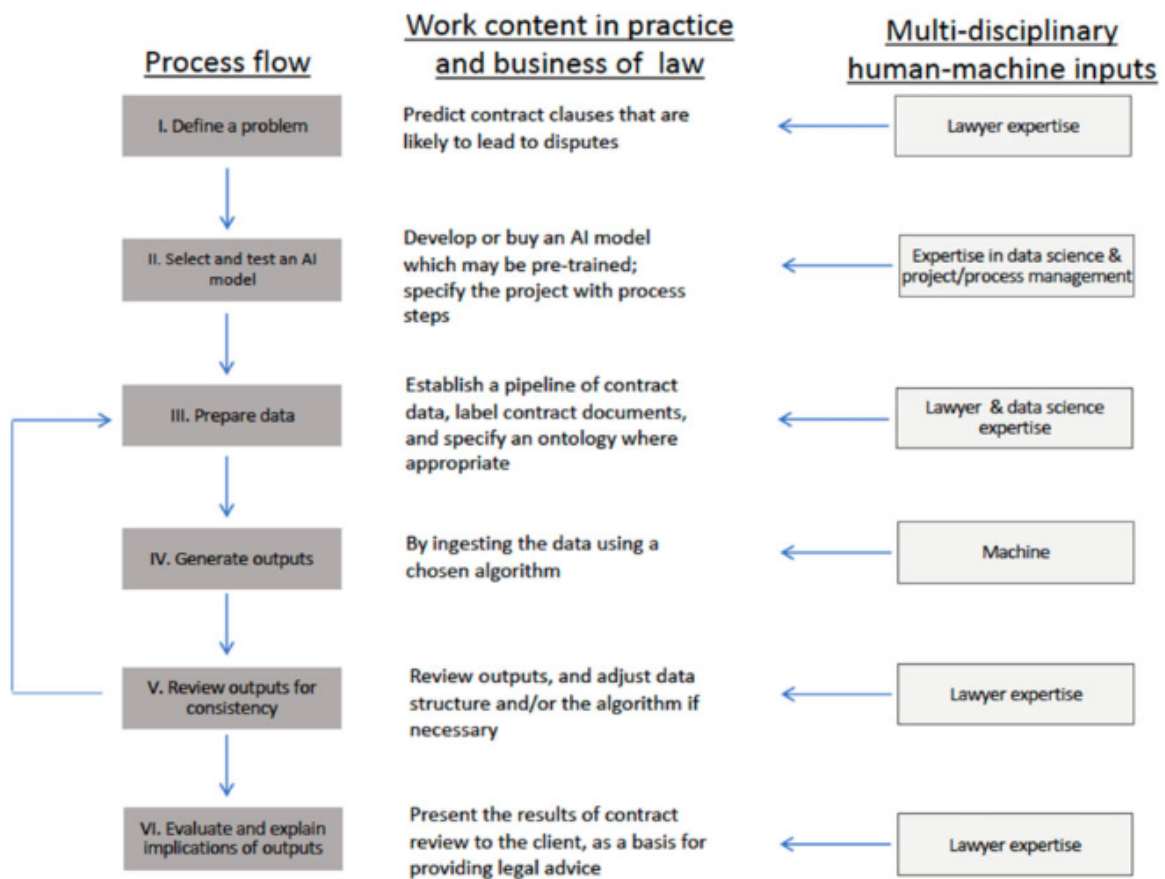


Figure 1: Process steps in AI use and step-specific requirements in human resources

Source: Armour and Sako (2020)

As the above process illustrates, five of the six steps require human expertise, drawn from lawyers, project/process managers and data scientists. This approach to implementing AI resonated with my own professional observations, through which I had seen increasing numbers of lawyers leave traditional fee-earning roles, to apply their legal knowledge as part of a multidisciplinary team responsible for digitising the firm’s legal processes.

Armour and Sako (2020) ultimately used their analysis to generate a typology of AI-enabled law firm business models, thus returning to the macro-level aspects of AI research discussed previously. What their process map highlighted to me, however, was the critical importance of the interactions between different professionals, and professionals and technology, for law firms moving to AI-enabled business models. While inter-profession relationships are attracting increased attention (Noordegraaf, 2020), the lack of research focusing on how

professionals interact with AI (at the time my initial literature review took place), encouraged me to turn to the more general literature on human-computer interactions from the field of Information Systems, which is discussed in section 2.3.

Identity theory offers another important theoretical perspective for understanding the behaviour of professionals. The concept of professional identity relates to how professionals define their self-image, in terms of their work role (Chreim, Williams and Hinings, 2007). Professional identity can, therefore, be seen to reflect the answers to two questions '*who am I as a member of this profession?*' and '*what do I do?*' (Nelson and Irwin, 2014). In most circumstances, individuals seek to maintain a coherent and positive self-image, which leads them to interpret workplace experiences in ways that are consistent with their professional identity. The tight connection between identity and work also tends to promote stability, as for one to change the other needs to change too. This can result in professionals resisting changes to their working practices that are imposed upon them (Chen and Reay, 2020). However, the desire to maintain consistency can also mean that professional identities are regularly redefined in light of new events that highlight inconsistencies between existing identities and actual work experiences (Nelson and Irwin, 2014). Different explanations have been offered as to *how* professional identities are constructed. Drawing on the distinction made by Goto (2021), I now outline two different perspectives of identity construction.

Within the professional services field, identity construction has more commonly been explained through reference to institutional logics. In this perspective, logics are understood to regulate the identity and behaviour of individuals; meaning changes in logics can be expected to cause shifts in identity and behaviour. Historically, the training and organisational norms that professional associations promote amongst their members are regarded as important in shaping professional identity (Anteby, Chan and DiBenigno, 2016). A competing set of logics highlighted by critical management scholars, can be found in management discourses that promote identity construction that is aligned to management objectives. In the context of professionals, Alvesson and Willmott (2002) argue these logics have sought to shift the meaning of professionalism, so that individuals prioritise the interests of their firm, ahead of their wider profession.

Further developing the logics perspective of identity, Bévort and Suddaby (2016) offer the concept of identity scripts; this ascribes individuals a higher degree of agency in identity construction. They argue self-authored identity scripts better explain how professionals interpret the competing institutional logics they are offered, and how they choose to enact them through their work. This helps explain why professionals working in a shared context and who are exposed to the same set of logics, typically show variation in the identities they construct and their workplace behaviour. This suggests professionals revise their identity scripts, in response to the changed institutional logics available to them following the introduction of a new technology, rather than as a direct response to the technology itself.

An alternative perspective of identity construction, which is more commonly seen within IS research, regards technology as an environmental condition, meaning that an individual's identity shifts in direct response to how they experience and interpret technology (Korica and Molloy, 2010). Depending on the evaluation they make, individuals can conclude that the technology is an opportunity or a threat to their self-image (Petriglieri, 2011).

In research amongst medical professionals who were required to use AI software at work, Jussupow *et al.*, (2018), identified five dimensions of threat linked to the introduction of new technologies – threats to *expertise, status, autonomy, influence* and *professional values*. Threats to expertise and status are categorised as an *individual-directed threat* as these are experienced by the individual as a threat to their own professional position within a given context. In contrast *group-directed threats* (autonomy, influence and professional values) are understood to devalue how professionals understand their professional role. Where an identity threat is determined, individuals are likely to exhibit one of two responses. *Identity protection* seeks to resolve the issue by devaluing the importance of the threat, or the identity that is being threatened. This can manifest in behaviour, such as resistance to change and rejection of the technology. An alternative response is *identity change*, through which the individual adapts their professional identity to make it consistent with the experience. In such cases, enacting the new identity can lead professionals to exhibit new behaviours, such as making a new technology part of their professional practice.

Relating the above to the phenomenon of AI-enabled legal services I anticipated that identity theory could help me explain variations in the behaviour of legal professionals

towards AI-enabled legal services. Drawing on both perspectives of identity construction in my analysis would allow me to evaluate whether the observed behaviour of professionals better reflected their direct experience of AI, or their interpretation of the institutional logics available to them. More specifically, Jussupow *et al.*'s (2018) five dimensions of identity threat could provide a useful framework for explaining why some lawyers choose to adopt AI-enabled legal services, while others choose not to, as there are many similarities in how medicine and law are structured as professions e.g. in terms of education and training, and the recent impact of AI on professional practice. I, therefore, determined that professional identity would be an important topic to explore, during the process of data generation.

Emerging professions and new models of professionalism

Models of professionalism seek to explain how professions define and organise themselves at a macro level, to deliver and protect their expertise. The legal sectors of Anglo-Saxon countries, including the UK, are regarded as archetypes of collegial professionalism, meaning they are granted a high-level of autonomy in both defining and regulating the scope of their work (Ackroyd and Muzio, 2007). In contrast to this, more recent knowledge-intensive occupations, such as management consultancy, have developed their own novel models of professionalism that are better able to manage the professionals and organisations that constitute their profession (Brock, Leblebici and Muzio, 2014). This distinction has previously been presented as a contest between models of *professionalism* (the treatment of complex cases by professionals for clients) and *managerialism* (the efficient delivery of products and services that meet the needs of customers).

A number of external pressures have also been identified as impacting contemporary models of professionalism, these include: financial-economic, socio-economic and technological (Noordegraaf, 2020). Financial-economic pressures can be seen through the rising costs of delivering professional services, and pressure from clients to reduce the amount they spend on professional advice. Together these have driven an organisational logic of efficiency, rather than the traditional professional logic of quality. Socio-cultural pressures are reflected in the changing demands of clients and increasing complexity of the cases professionals are asked to treat. In some instances, this complexity can mean no

single professional is capable of solving the problem as presented, meaning there models of professionalism that enable coordination and collaboration across jurisdictional boundaries. Technological pressures have seen the digitalisation of professions, with professional decision-making and treatment increasingly supported by tools such as AI and robotic tools.

These pressures have helped move the debate beyond the professional-management dichotomy, with concepts such as hybrid professionalism (Noordegraaf, 2015) arguably better reflecting the intertwined nature of professionalism and managerialism in many organisations from the 1990s onwards. In hybrid model of professionalism, professionals are still understood to be focused on the treatment of cases, but within the context of modernised and well-managed organisational structures, such as the professional services firm (Empson et al, 2015), that provide an environment that promotes high levels of coordination amongst professionals.

While models of professionalism ultimately reflect a range of economic and social factors including technology, technological change is recognised to have impacted both the nature and volume of cases that clients consult professionals about; and the ways in which professionals diagnose and treat these cases. To better understand the impact of technology on professional practice requires a shift from the macro to the micro level. In their analysis of the professions, Abbott (1988) conceptualises professional practice as comprising of three different acts – *diagnosis*, *inference* and *treatment* (see Figure 2). Diagnosis involves analysing the case as presented by a client, this means taking information from/about the client (evidence) and framing it against the professional's own knowledge base. The act of treatment requires professionals to use the techniques and tools available to them to deliver the prescription that they have identified to be most suitable to resolve the client's case. When a case is well-understood it is relatively straight-forward to identify the relationship between the inputs identified during diagnosis, to the expected outcome of administering a specific treatment. In such instances the process of diagnosis and treatment can be almost automatic and rules-based, involving limited risk and requiring minimal professional judgement. However, when the connection between diagnosis and treatment is obscure, professionals need to engage in the act of inference, drawing upon their tacit knowledge and expertise, often in creative ways, to arrive at a valid conclusion to

the case. Through understanding professional practice as comprising these three acts, it becomes possible to see how technology has the potential to influence the work of professionals in different ways, depending upon which act(s) the technology interacts with.




Figure 2: The process of knowledge work (adapted from Abbott, 1988)

Source: Harvey, Heineke and Lewis, 2016 (2016)

For individual professionals, technological change has, therefore, led to their work being reconfigured in terms of new organisational processes and patterns of interaction. The precise outcome of such changes is difficult to predict (Köktener and Tunçalp, 2021), and can reflect both the materiality of the technology and jurisdictional disputes between different groups of workers (Barrett *et al.*, 2012). Köktener and Tunçalp's (2021) analysis of role reconfiguration amongst auditors, using Abbott's (1988) three acts of professional practice framework, suggests that reconfiguration is more likely to take place during the acts of diagnosis and treatment, but that the encroachment of technology (or other types of professional) into the act of inference will be resisted. This means for accommodation to be reached the act of inference is likely to continue to be conducted solely by professionals with the relevant subject-matter expertise.

In the legal sector, the growing number of jobs in multidisciplinary teams responsible for AI-enabled legal services, and the pay premium offered to lawyers with digital skills provides some empirical evidence that reconfiguration is taking place (Sako, Qian and Attolini, 2022). What has not yet been demonstrated, is *where* the reconfiguration has taken place; and *who/what* is undertaking the work. This led me to conclude that there was an opportunity

to use my research to better understand where AI-enabled legal services were reconfiguring roles, and whether the pattern reflected that seen in other professions. I, therefore, recognised that during the data generation phase of my research I would need to identify the individuals and technologies involved at different stages of the process used to deliver AI-enabled legal services, and that this would require participants to be drawn from both traditional fee-earning roles and multidisciplinary teams.

Researcher reflections

My over-arching reflection from reading AI-related research conducted within the professional services field, was that researchers seemed more interested in understanding the macro-level impact of AI on professions and organisations, with far less research focusing on the behaviour of professionals towards AI. While I was already predisposed to focus on the phenomenon at the level of the individual, my review of the literature gave me confidence that my research could make a theoretical contribution by focusing on micro-organisational issues relating to AI. I was also struck by the relatively abstract conceptualisation of AI as an exogenous force that was used in much of the research. This left me feeling unclear about how AI was being distinguished from related phenomena, such as digitalisation. I resolved to focus my research on the more specific phenomenon of AI-enabled legal services, in anticipation that this would help to make AI technology visible, thus allowing me to scrutinise it and better conceptualise it as a phenomenon. My final reflection was that several of the useful concepts and theories I had identified in the literature stemmed from academic disciplines I had limited knowledge of. This led me to conclude that my own research would benefit from being phenomenon-led, rather than theory-led; and that in seeking to understand AI-enabled legal services I would need to conduct my research using a theoretical standpoint that would enable me to employ an interdisciplinary approach to analysis (discussed in Chapter 3).

2.3 Review of the Information Systems literature

In seeking to identify the theoretical foundations of information systems (IS) research, Lim *et al.* (2013) identified five research streams, based on an analysis of the leading IS journals

from 1998-2006. Of these five streams, I judged that two mapped relatively well to research streams identified by Brock, Leblebici and Muzio (2014) within the professional services field. *IT and Organisation* research focused on the role of information technology at the organisational level, for example, through its contribution to strategy; effect on internal processes; and as a general enabler of organisational performance. Within this stream, over two thirds of the theories cited in research were drawn from the disciplines of strategy, organisational science and economics (Lim *et al.*, 2013). The subject matter of this stream, therefore, seemed to overlap with the *Organisational models and structures* stream identified by (Brock, Leblebici and Muzio, 2014).

Of greater relevance to my own research was the *IT and Individuals* stream, which focused on human-computer interactions. This stream was underpinned by theories from the disciplines of psychology, information systems and sociology, which together accounted for over 80% of the theoretical base in published research. The dominant theory, which featured in 19% of research, was Davis' (1989) technology acceptance model (TAM), which seeks to explain how users come to accept and use a technology. The *IT and Individuals* stream, therefore, seemed to consider topics also found within the *Micro-organisational issues* stream (Brock, Leblebici and Muzio, 2014), specifically how individuals respond to the introduction of new technologies. This led me to centre my theoretical literature review on human-computer interaction research (from now on referred to as user-system interaction research) that had a technological focus on AI.

The tool perspective of technology

The tool perspective of technology treats technology as a human-created artifact, designed to serve a human-defined purpose (Lee, Thomas and Baskerville, 2015). In process terms, human *users* provide inputs through an artificial *interface*, to a technological *artifact* that in response generates an *output*. In the context of my research, we can think of a lawyer (the user) dictating into a computer microphone (the interface) that transforms the sound of their voice into an electrical signal; their computer (the artifact) then processes this input using its software, transcribing it into text (the output) that can be viewed on a screen.

The IS field has utilised the tool perspective of technology for several decades, which has led it to become an important assumption upon which several theories within the *IT and Individuals* stream of research have been built (Schuetz and Venkatesh, 2020). For example, the Technological Acceptance Model (Davis, 1989) and its iterations (Gefen, Karahanna, and Straub, 2003; Venkatesh *et al.*, 2003) predict technology use is influenced by the technology's *perceived usefulness* and *perceived ease of use*. This theory only makes sense because it is assumed there is a human user and a technological tool for them to use.

How AI challenges the tool perspective of technology

More recent research has questioned the suitability of the tool perspective for theorising about novel technologies, which do not behave in the passive manner of a tool. Ågerfalk *et al.* (2022), propose that in trying to understand how AI resembles (and differs) from other technologies there is a need to reconceptualise the notion of agency as being necessarily human, and instead recognise that certain technologies can be better understood as digital agents. Indeed, in certain circumstances, it is argued, it can be more accurate to consider humans as artifacts, which are shaped and used by the technological systems that surround them. Information systems that employ AI to make decisions, such as algorithmic trading systems within the financial services sector, are one example of this (Demetis and Lee, 2018).

This means that while the IS field has developed substantial knowledge about information systems and how humans interact with them, the extent to which this knowledge translates to theorising about AI remains an open question. There is, therefore, a pressing need for empirical research that can reduce the uncertainty surrounding how to best manage AI, and maximise the likelihood that it has a positive impact (Berente *et al.*, 2019).

Before using an existing theoretical framework to explore AI-enabled legal services, it is therefore, important to consider which field-level assumptions the phenomenon is most likely to contest; a difficult task given such assumptions are often absent from discussion (Sandberg and Alvesson, 2011). Helpfully, Schuetz and Venkatesh (2020) identified five user-system interaction assumptions stemming from the tool perspective, which could be

undermined by Cognitive Computer Systems (CCS), a type of AI that is designed to simulate human thought (see Table 5).

Table 5: How Cognitive Computer Systems challenge IS assumptions

Element	IS Assumption	Assumption label	Scope	Challenged by CCS characteristic
User	1. Humans are users	Unliteral relationship	Field assumption (e.g., system adoption, privacy)	Interactive
Input	2. The developer defines the inputs	Ignorance of environment	Field assumption (e.g., system adoption, privacy, communication)	Adaptive, interactive
Computation	3. IT artifact use leads to consistent outcomes	Functional consistency	Field assumption (e.g., IS success, IT governance, IS development)	Adaptive
Output	4. The way the tool derives its outcomes is comprehensible and can be verified	Functional transparency	Field assumption (e.g., recommendation systems, IS development)	Contextual, adaptive
Interface	5. There is an artificial interface	Awareness of use	Field assumption (e.g., privacy, service science)	Interactive, iterative and stateful, and context aware

Source: Schuetz and Venkatesh (2020)

While it is true that not all AI-enabled legal services can be considered CCS (although some would meet this definition), the assumptions highlighted by Schuetz and Venkatesh (2020) provided me with a useful framework to think about the ways in which AI-enabled legal services can differ from other technologies used by legal professionals, and what IS assumptions they challenge. The embryonic nature of AI-enabled legal services as a phenomenon, meant there was a lack of published research I could reflect on. Instead, I used the framework to help me reflect on my own observations of AI use in law firms and analyse the technical product information published by Legaltech service providers. The extent to which I determined AI-enabled legal services have challenged the five assumptions are detailed below. My analysis is illustrated through reference to existing use cases of AI in law firms.

1. AI-enabled legal services are tools used by humans.

IS theories typically assume a user-artifact relationship between human and technology (Norman and Draper, 1986), however, the parameters of certain AI systems give them the ability to influence the behaviour of humans, in pursuit of the system's (human-designed)

goals (Demetis and Lee, 2018). In these instances, it is more accurate to regard AI as an active agent, than a passive tool. For example, a law firm might want to improve the ability of their lawyers to accurately predict the outcome of a case. To achieve this goal, they deploy an AI-enabled legal service which transforms the firm's historic client data into a smart report. The report provides a summary of data insights, highlighting the probability of different case outcomes, and a recommended course of action to follow. The report is emailed to the relevant lawyer automatically when they open a new case file for a client, without the need for them to request it. The information in the report forms the basis of the case strategy advice the lawyer gives their client.

While not all AI-enabled legal services will break the assumption of a unilateral relationship between the legal professional and the technology they use, in some instances this assumption can already be seen to have been broken. For example, Solomonic software provides a significance grade when presenting data, which provides lawyers with a defensible context for their legal advice (Solomonic, no date). AI can, therefore, be seen to have the potential to affect the way in which professionals perform all three of Abbott's (1988) professional acts.

2. AI-enabled legal services are ignorant of their environment.

IS theories generally assume the technological artifact operates in isolation from (and in ignorance of) its surrounding environment, with system inputs typically controlled through an interface by a human user or another computer. This assumption no longer holds when AI has been designed to directly access data itself, for example, through sensors or a connection to the internet; it can then use this data to pursue its goals, without direct human supervision (Demetis and Lee, 2018). To continue the previous example, to further enhance the ability of lawyers to predict cases, the firm deploys an AI-enabled legal service underpinned by software which tracks UK High Court litigation data on an ongoing basis. This allows external real-time data (e.g. the rulings of individual judges or the success rates of lawyers at other law firms) to be analysed and combined with the firm's own historic data (which was inputted into the system by the firm's IT developers). The resulting AI-generated report contains insights based on a volume of data that no lawyer at the firm would have been capable of accessing and analysing for themselves.

While many law firms choose not to develop AI-enabled legal services that can independently access their environment, some Legaltech already possesses this ability e.g. Solomonic can access UK High Court litigation data (Solomonic, no date). This suggests AI-enabled legal services have already challenged this assumption.

3. AI-enabled legal services are functionally consistent.

A widespread assumption is that system functionality is determined by the developers and purchasers of information systems, rather than the system itself. The adaptive nature of some AI models means that while their intended goal remains fixed, the functionality they employ in pursuit of this goal can change. This happens when their internal statistical models are refined in response to new data and system feedback, without the need for direct human intervention. For example, a law firm might want to make it easier for lawyers to identify any non-standard clauses in a data set of several thousand documents. The firm implements an AI-enabled legal service that uses natural language processing software to identify any clauses that do not perfectly match a pre-defined template, specified by a legal subject matter-expert at the firm. These clauses are then flagged automatically to a lawyer, who decides whether any action needs to be taken. The lawyer then provides feedback to the AI software indicating whether the flagged document was pertinent or not. Based on this feedback, the software independently updates its algorithm so that it is better able to identify relevant documents when a similar query is made in the future.

Many AI-enabled legal services utilise machine learning techniques, which allow the AI software to develop more accurate algorithms over time, for example, Kira Systems contract review software (Kira Systems, no date). This means that when used for an extended period of time, some AI-enabled legal services cannot be assumed to be functionally consistent.

4. AI-enabled legal services are functionally transparent.

IS research also assumes that systems behave objectively, according to their programming, whereas the human users of these systems are subject to cognitive bias. The assumption that a system will always make correct calculations, based on the inputs it is given, is an important component of theories of trust in technology (Lankton, McKnight and Thatcher,

2014). The adaptive nature of some AI software challenges this assumption, as it permits the generation of outputs based on inputs or processes the user is either unaware of, or unable to understand; this means AI systems can lack transparency. The importance of this can be seen by returning to the example of a law firm with the goal of improving the ability of their lawyers to predict case outcomes. Unlike a lawyer who might generate their case forecast using a mental model, which combines the legal facts of the case with their previous experience of similar cases; the firm introduces an AI algorithm that analyses a much larger data set containing several thousand data points, each of which is individually weighted in the AI's decision-making model. Assuming the data is accurate and relevant, the accuracy of the AI model is likely to be much greater than the mental model of the legal professional, but the overall complexity of the model may mean it is impossible for a lawyer (or their client) to understand how the system arrived at its answer.

To date, there are no regulations in place requiring law firms to use algorithms that meet a specific standard of transparency. Given the aforementioned use of complex machine learning models in some Legaltech software, it is reasonable to assume that the assumption of algorithmic transparency has already been broken by AI-enabled legal services.

5. Awareness of using AI-enabled legal services

It is assumed that users are aware of when they are interacting with an information system, as the interaction takes place through an artificial interface e.g. a keyboard. This means the user's behaviour during such interactions can be deemed intentional. The ability of certain AI systems to exhibit human-like behaviours means users can have difficulty recognising whether they are interacting with a human or artificial agent, if the interaction between the two is mediated by technology. For example, a law firm might wish to reduce the average length of time it takes to respond to client queries. The firm introduces a chatbot that clients access via the firm's website. Utilising generative AI and natural language processing, the chatbot is designed to provide written responses to simple legal queries and collect relevant information from clients. The responses are drafted by the AI chatbot in a professional yet conversational tone that follows the firm's own copywriting style guide; making the responses difficult to distinguish from those drafted by the firm's lawyers. The

clients appreciate the speed of response, but some incorrectly assume they were interacting with a human.

Chatbot use is not yet widespread amongst law firms. Where it has been introduced, chatbots like Harvey (Harvey, no date) have been made available to legal professionals employed by the firm (rather than to clients), who it can be assumed are aware they are using a chatbot. Of the five assumptions identified by Schuetz and Venkatesh (2020) this is the one that is least likely to have been undermined by AI-enabled legal services.

Researcher reflections

Schuetz and Venkatesh (2020) identified a number of opportunities to conduct interesting research and develop theory (Davis, 1971) when each of the five assumptions is relaxed. However, their suggestions are context-neutral, and not all are of immediate relevance to AI-enabled legal services. I, therefore, reflected on the ways in which AI-enabled legal services might necessitate the revision of IS theories, based on a combination of my professional experience and my review of the professional services literature. I now explain the impact this had on my research, highlighting relevant theories I decided to consider in greater detail, and the methodological implications relating to data generation.

From a technical perspective, I anticipated it would be relatively straight-forward for a subject-matter expert with knowledge of AI, to objectively determine whether the technical specifications of the AI system deployed in their organisation conformed to the tool perspective or challenged one or more of the five assumptions. In firms where multiple AI-enabled legal services were in use, each service would need to be assessed separately according to the technical parameters of the AI system that underpinned it. Many of the legal professionals who utilise AI-enabled legal services as part of their professional practice, do not possess detailed knowledge of AI. While this meant they would be unable to make an accurate technical assessment of its capabilities, they could still explain whether or not they perceived it to be a tool, based on their personal experiences. This led me to conclude that I would need to generate data with a range of participants, some of whom would need to possess expert knowledge of AI. It would also be necessary to use different types of

questioning, designed to reflect the kind of knowledge that was of interest (Bogner, Littig and Menz, 2009).

An important goal of IS theory is to explain levels of technology use. The user-tool perspective is an important component of these theories, which meant it was important to find out whether a unilateral relationship accurately described how legal professionals interacted with AI-enabled legal services. Answering this question would also help explain whether the use of AI by legal professionals followed a similar path to the acceptance of other technologies; and the extent to which the outputs of AI-enabled legal services influenced the behaviour of legal professionals. I, therefore, recognised that I would need to learn about theories of user acceptance, prior to generating data that would allow me to understand how legal professionals perceived their relationship with AI-enabled legal services; and whether this was a different kind of relationship to what they experienced when using other types of technology.

I also surmised that AI-enabled legal services with environmental awareness would have the potential to augment the decision-making of legal professionals. Where legal professionals recognised this benefit, the perceived value of AI-enabled legal services would be significant (AI might even be regarded as a necessity), leading to increased levels of AI use. I, therefore, resolved to generate data that indicated how legal professionals understood AI-enabled legal services to have impacted their professional practice, and whether their work was dependent upon the use of AI technology.

Informal discussions within my professional network, which helped inspire my research, suggested that not all legal professionals trust AI technology, and that this was hampering levels of adoption within the sector. This observation reflects a wider belief that lawyers display lower levels of trust than the general population (Richard, 2002). This made me aware that I would need to familiarise myself with the literature relating to trust in technology (Mcknight *et al.*, 2011). Exploring how trust in AI-enabled legal services developed amongst legal professionals, would generate empirical data that could be used to assess the extent to which existing theories accurately described trust development amongst legal professionals. It would also allow me to see whether functional consistency and functional transparency were important to trust development in AI-enabled legal

services. To achieve this, I needed to generate data that would indicate whether legal professionals trusted AI-legal services (or not), and what factors underpinned this trust (or prevented it from developing). Given that trust in technology usually develops over time (Schaefer *et al.*, 2016), I recognised that it would be important to ask participants about the level of trust they experienced throughout the time in which they were using AI-enabled legal services, and not just focus on their level of trust at the moment the data was generated.

2.4 Trust and acceptance of artificial intelligence

The research opportunities identified above, led me to investigate the theoretical literature relating to user acceptance of technology and trust in technology. Familiarising myself with this literature prior to commencing data generation and analysis helped me evaluate the extent to which these theories provided a suitable framework for investigating AI-enabled legal services and sensitised me to how the relaxing of the aforementioned assumptions might limit their application.

Underpinning theories of user acceptance and trust are two broader theories of human behaviour that have significantly influenced IS research – the Theory of Reasoned Action (Fishbein and Ajzen, 1975) and Theory of Planned Behaviour (Ajzen, 1985). Their importance reflects a focus within the IS field on phenomena that are difficult to observe and measure, for example user-system interactions. This has led researchers to use proxy measures, rather than measuring the behaviour itself (Tao, 2009; Granić, 2023). For example, a researcher interested in measuring the use of AI-enabled legal services, might focus on attitude towards AI, or reported intention of a legal professional to use AI. This could be measured in a relatively straight-forward manner using a standardised scale administered by survey. In contrast, measuring AI use through monitoring the actual behaviour of the legal professional would require detailed observations to be made over a period of time, or access to electronic user data, assuming the AI system can provide this.

The Theory of Reasoned Action (TRA) provides a framework that justifies researchers drawing inferences between an individual's intentions and their actual behaviour. The theory suggests that prior to *behaviour* taking place, an individual first forms an *intention* to

undertake the behaviour, based on a rational thought process. Intention reflects both *attitude* (the beliefs an individual has towards a behaviour) and *subjective norms* (the perceptions of important others about the behaviour). To understand an individual's intentions both the direction (positive or negative) and weighting of attitudes and norms must be considered (see Figure 3). It is also important to note that intention does not perfectly predict behaviour, as attitude and norms are dynamic; meaning intentions can change before the anticipated behaviour occurs (Conner and Norman, 2022). In the context of researching AI-enabled legal services, the intention of a legal professional to use AI would be seen to reflect a combination of their own attitude towards AI-enabled legal services, and the normative attitudes held by important colleagues and clients. In a context such as the UK legal profession, which has long-established norms that shape what is considered professional conduct, we might expect subjective norms to play a more significant role in determining intentions and behaviour than in other organisational contexts.




Figure 3 - Theory of Reasoned Action (Fishbein and Ajzen, 1975)

The limited empirical success of the TRA in predicting behaviour led Ajzen to introduce an additional construct, *perceived behavioural control*, to better explain the formation of behavioural intention; this revised model was called the Theory of Planned Behaviour (Ajzen, 1985). Perceived behavioural control draws upon self-efficacy theory (Bandura, 1977), and refers to the ease (or difficulty) an individual believes they would have in performing a behaviour. The model predicts a positive relationship, whereby the greater the level of perceived behavioural control, the greater the behavioural intention. This means when an individual's control over their behaviour is constrained, for example as a result of organisational rules and practices, or the affordances of the technology, their intention to behave in the desired way is reduced, even though their attitude towards that behaviour may remain unchanged (see Figure 4). In a law firm context, a firm-wide policy that makes the use of AI-enabled legal services mandatory, would reduce the perceived

behavioural control of legal professionals to decide for themselves whether to use AI or not, irrespective of whether they held a positive or negative attitude towards AI-enabled legal services. In such circumstances, an individual with a negative personal attitude towards the use of AI may still intend to use the technology because they lack the perceived behavioural control necessary to perform their work without using the firm's AI systems.




Figure 4 - Theory of Planned Behaviour (Ajzen, 1985)

While the Theory of Planned Behaviour (TPB) has been shown to be a better predictive model of behaviour than the TRA, it is not without its critics. It has, for example, been indicated that TPB fails to account for the impact of emotions in decision making and that the relationships between the factors identified by Ajzen (1985) may in fact be more complex than suggested. For example, research conducted by Sussman and Gifford (2019) demonstrated a bi-directional relationship between intentions and its antecedents, that was not predicted by the model. This helps explain why meta-reviews of the theory have found significant variation in the strength of the relationship between intention and behaviour (McEachan *et al.*, 2011) which in part helps shed light on the fact that the TBP accounts for less than a third of the variances in observed behaviour (Armitage and Conner, 2001)

In spite of their limitations, these two theories provide a useful set of concepts for thinking about the behaviour of legal professionals in relation to AI-enabled legal services. While the models do not perfectly predict behaviour, I anticipated it could still be valuable to draw upon elements of the models when analysing my data to explain patterns of behaviour amongst legal professionals using AI. I, therefore, decided that it would be important to generate data about the attitude legal professionals had formed towards AI-enabled legal services; how other parties viewed the use of AI; and the extent to which they believed they

had control over their interactions with AI. I now discuss relevant theories of user acceptance and trust that have been developed from these theoretical underpinnings.

Theories of user acceptance of technology

The concept of technology adoption refers to the acceptance, integration, and embracement of a new technology, with technology acceptance regarded as the first step of technology adoption (Granić, 2023). Technology acceptance is a major challenge faced by those who seek to promote technology use, which explains the importance of user acceptance theories within the IS field. This has led to the development of several theories that seek to explain the differences that can emerge between the predicted and observed behaviour of system users. Drawing on research from IS, psychology and sociology these theories have enjoyed varying levels of success in predicting differences in the intention of individuals to use a technology. In their review of user acceptance models Venkatesh *et al.* (2003) identify eight competing models, which together are comprised of over 30 different constructs that it is suggested are predictors of the *intention to use*. Of the models featured in the review, the Technology Acceptance Model (Davis, 1989; Venkatesh and Davis, 1996) is the most frequently cited theoretical framework within the *IT and Individuals* research stream (Lim *et al.*, 2013). I, therefore, chose this as my starting point for familiarising myself with the user acceptance literature.

The Technology Acceptance Model (TAM) states that an individual's *intention to use* a system is predictive of their *usage behaviour*, and that *intention to use* reflects their beliefs about the system's *perceived usefulness* and *perceived ease of use* (see Figure 5). *Perceived usefulness* is understood to mean the extent to which an individual believes a system will help them to achieve their personal goals. In the context of AI-enabled legal services this would typically mean the extent to which a legal professional believes using AI would help them to hit their individual performance targets and realise their career goals. *Perceived ease of use* reflects the degree of effort the individual believes is required on their part in using the system to achieve their goals. For example, a legal professional's belief about the amount of effort required to complete a task using AI, could be contrasted with completing the same task without the use of AI.

In empirical testing of the original TAM (Davis, 1989) it was found that both *perceived usefulness* and *perceived ease of use* had a direct relationship with *intention to use*, without *attitude towards using* being required to mediate the relationship as was originally thought. This led to a revised and more parsimonious model of TAM (Venkatesh and Davis, 1996).

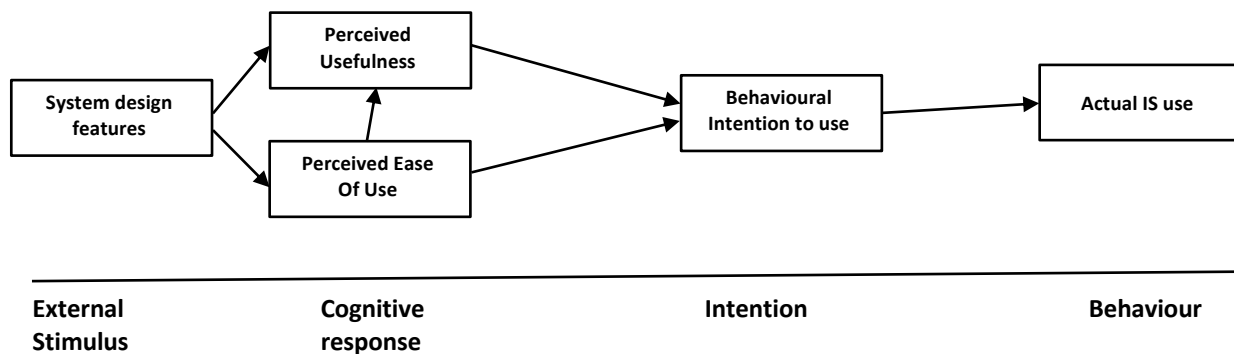


Figure 5 - Technology Acceptance Model (Venkatesh and Davis, 1996)

This was then elaborated upon in TAM2 (Venkatesh and Davis, 2000), which sought to explain what factors led an individual to perceive a system to be useful. TAM2 included a range of different antecedent factors, including: the perceived *relevance* of the system to the individual's job role; the perceived *quality* of the output produced by the system; and the first-hand *experience* of the individual once they had started to use the system.

Researcher reflections

The user acceptance literature was helpful to me in identifying the ways in which the characteristics of AI-enabled legal services as a system, might influence the likelihood of legal professionals using the system. By treating technology as an independent variable these theories highlight the direct role technology can play in influencing behaviour, thus making the technological artifact more visible to researchers. By introducing specific technology-related constructs, such as *perceived usefulness*, these theories gave me a range of factors to analyse my findings, in addition to those identified in the professional services literature.

However, the range of different theories and variables identified in this research stream (Venkatesh *et al.*, 2003) indicated that user acceptance theory has struggled to fully explain empirical findings. While this may in part reflect theory development lagging behind the

rapid rate of technological change in organisations (Orlikowski and Scott, 2008), it also indicated to me that the complexity of user acceptance, was only likely to be explained through reference to multiple theories. This more inclusive theoretical approach is supported by Jensen and Aanestad (2007) who argue that TAM's cognitive theoretical focus, fails to properly explain the sources of the perceptions that guide acceptance behaviour; and overlooks the complex relationships that can develop between user, system and its organisational context. These broader considerations can instead be better understood through reference to theories such as technological frame theory (Orlikowski and Gash, 1994). The limitations of IS theories of user acceptance are further exacerbated by the challenges AI poses to the tool perspective, which underpins models of user acceptance. Specifically, the assumptions of functional consistency and functional transparency, which allow user-acceptance models to treat technology as a constant variable, may not be justified, making variance research using quantitative methods problematic (Schuetz and Venkatesh, 2020).

I therefore concluded, that while it was important that my research considered the direct impact AI-enabled legal services had on the behaviour of legal professionals, this was likely to be just one of several factors that influenced levels of AI use. As each theoretical framework would only provide a partial explanation, I decided to adopt a theoretical stance capable of generating a localised causal explanation of user acceptance. Methodologically this meant selecting an approach to data generation and analysis that would allow me to consider both the role of AI as an artifact, and the social context it was embedded within, when explaining the use of AI-enabled legal services by legal professionals.

Theories of trust

Trust is studied across disciplines, including psychology, sociology, and economics; perhaps unsurprisingly, this has led to different understandings of trust. In psychology, trust research typically focuses upon personality development, interpersonal relationships and how individuals relate to the wider world. In economics, trust is seen to help facilitate market efficiency by reducing transaction costs between parties. In sociology, trust is studied at both the micro level through the behaviour of social actors, while at the macro level trust is theorised to be an important part of modern and post-modern society. In the

field of information systems, existing theory suggests that trust is significant to a human's willingness to interact with technology (Orlikowski and Scott, 2008) and that without trust other predictors of technology use become less relevant (Liu and Goodhue, 2012).

In spite of this diverse use of the term, Rousseau et al (1998) demonstrate that at their core, different definitions of trust reference the positive expectations and vulnerability of an individual, and their willingness to depend on another party, because of that party's attributes. Hoff and Bashir (2015) develop this further, identifying three common elements of trust:

1. There is a *truster* to give trust, a *trustee* to accept trust, and something must be at stake.
2. The trustee is incentivised to perform the task. For example, by a benevolent internal desire to help, or the promise of financial reward.
3. There is uncertainty or risk, which means the trustee could fail to perform the task.

Together these outline the notion that in situations of uncertainty, trust is a necessary condition in order for cooperative exchange to take place. This is reflected in Gefen's (2004) definition of trust as,

"The belief that others upon whom one depends, yet has little control over, will not take advantage of the situation by behaving in an opportunistic manner but, rather, will fulfil their expected commitments by behaving ethically, dependably and fairly, especially under conditions involving risk and potential loss." (p.264)

In the context of my research, this definition captures well the relationships of trust that develop between legal professionals who work interdependently as peers of one another, and whose behaviour is difficult to measure directly.

Trust in technology

The above definition conceives trust as a behaviour that can emerge between two or more humans (*human trust*). In the context of IS research this is sometimes an accurate reflection of the phenomenon being studied, for example when studying trust between different

human stakeholders in an IT project or members of a virtual team (Jarvenpaa, Knoll and Leidner, 1998). Human trust is also an appropriate concept when considering certain technology-mediated interactions, for example, between a user and a recommendation agent (Wang and Benbasat, 2008) or a user and an online intermediary (Du and Mao, 2018). In these cases, the trust relationship of interest is between the user and the human party that sits behind the website i.e. the product or service vendor, and not the technological artifact itself.

The ubiquitous use of technology across society means researchers are increasingly interested in trust between humans and technological artifacts, with a focus on the trust-related attributes of the artifact (Mcknight et al., 2011). This is variously referred to as *trust in technology* (for example, Li, Hess and Valacich, 2008; Lankton, McKnight and Thatcher, 2014) or *trust in automation* (for example, Hengstler, Enkel and Duelli, 2016; Schaefer et al., 2016). In human trust relationships the truster, gives their trust to the trustee, knowing the trustee may take advantage of this. However, trust in technology treats technological artifacts as lacking volition, meaning technologies cannot choose whether to behave as the trustor hopes and expects. Hence, Lee and See's (2004) more parsimonious definition of trust as, "*the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability*" (p.51). Thus, trust in technology refers to beliefs about how a technology will operate within a given work environment. For example, when a legal professional sends a confidential email to their client, they believe the technology will deliver the message promptly and securely. This distinction between human trust and trust in technology is important, because by focusing on trust in technology researchers can better determine what makes the technology itself trustworthy, irrespective of the humans and wider social structures that surround it (Mcknight et al., 2011). This allows evaluation of the extent to which the trusting beliefs a user has about the technological attributes of a system, influences their behaviour towards it.

Whether the focus is human trust or trust in technology, trust is important to understanding complex socio-technical phenomena, such as AI-enabled legal services. Trust's relevance to contexts characterised by risk and uncertainty, such as legal professionals providing advice to clients, means trust is an overarching concern for those responsible for the management

of individuals and technology (Lee and See, 2004). Indeed, trust’s role in promoting technology acceptance, through helping users overcome perceptions of risk and uncertainty, has seen trust incorporated into other established theories. Of relevance to this research, Gefen, Karahanna, and Straub (2003) developed the Trust-TAM model of user acceptance (Figure 6) by integrating trust as a further antecedent of *intended use*, alongside *perceived usefulness* and *perceived ease of use*, thus recognising the importance of trust to technology adoption alongside other established concepts, such as perceived usefulness and perceived ease of use. This model has had a significant impact on trust research and has been used to investigate the use of technology in a wide-variety of contexts including the use of e-commerce (Ha *et al.*, 2019), AI in medical diagnosis (Fan *et al.*, 2020) and mobile banking (Aldammagh, Abdeljawad and Obaid, 2021).

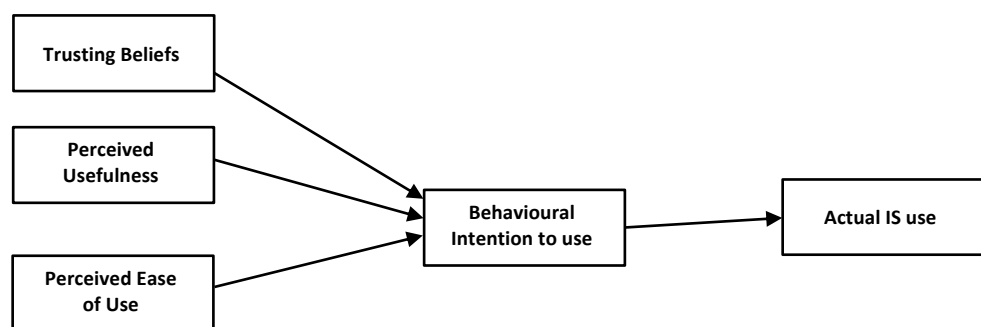


Figure 6 – Simplified Trust TAM model (adapted from Gefen, Karahanna, and Straub, 2003)

Conceptualising trust

The majority of IS trust research focuses on identifying *what* causes trust (its antecedents) and *what* consequences trust has in terms of behaviour. Such research conceptualises trust at the individual level, as a static, multi-dimensional, psychological state. Methodologically, this conceptualisation has led to a focus on cross-sectional, variance research using quantitative data, and a more limited focus on *how* trust develops (Rose and Schlichter, 2013). Drawing on the Theory of Reasoned Action (Fishbein and Ajzen, 1975), static models of trust indicate that *trusting beliefs*, which reflect the trustee’s attributes (sometimes referred to as their trustworthiness), influence an individual’s willingness to depend on the trustee (*trusting intention*). The greater the trusting intention, the more likely *trusting*

behaviour is to take place. In reality, however, trust can be more accurately understood as a dynamic concept that can change over time. This is not to suggest that trust necessarily develops gradually (Li *et al.*, 2008) instead, a distinction is typically made between *initial trust* and *experience-based trust*. What determines initial trust is believed to be different from trust with an already familiar party, as initially the truster has no direct experiences to draw upon when forming their initial trusting beliefs (Hoff and Bashir, 2015).

While a wide-range of factors have been proposed to explain the development of trust in information systems, meta-analyses of the literature have typically organised these factors into three over-arching categories – human-related, technology-related, and environmental (see Figure 7) – each of which has a range of dimensions within it (Hancock *et al.*, 2011; Schaefer *et al.*, 2016). All three of these need to be considered in order to understand why a legal professional may choose to trust an AI-enabled legal service or not.

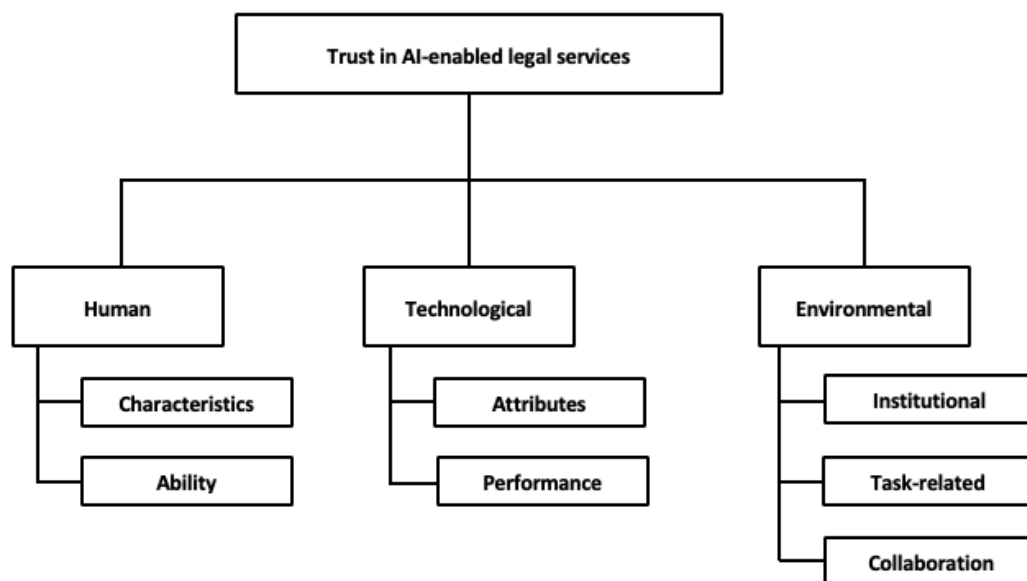


Figure 7: Factors of trust development in human-technology interactions

Source: Adapted from Hancock *et al.* (2011)

Research in *human trust* typically identifies *ability*, *benevolence* and *integrity* as the key trusting beliefs (Mayer, Davis and Schoorman, 1995), although *ability* is sometimes defined as *competence* (McKnight, Choudhury and Kacmar, 2002). When thinking about trust

between two legal professionals, it would be anticipated that trust is more likely when the individuals make a positive evaluation of one another's legal expertise and their individual character traits. Given the role of legal professionals in delivering AI-enabled legal services, I anticipated that human-related characteristics of both AI users and their colleagues, for example, those linked to personality and professional competence, would play a role in determining the level of trust shown in AI-enabled legal services.

In contrast, (Mcknight *et al.*, 2011) identify *functionality*, *helpfulness* and *reliability* as the corresponding set of beliefs typically adopted by *trust in technology* researchers, while Lee and Moray (cited in Hoff and Bashir, 2015) identify the *performance*, *process* and *purpose* of the technology as critical. While *reliability* and *performance* can be seen to be directly related to the quality of the output produced by a technology, *functionality* and *process* can be better understood as features of the technology, for example the design of its interface. Hence a legal professional deciding whether to use AI to complete a task, would be more likely to use AI if they were confident that it was technically capable of performing the task and that it would perform the task without error. The critical role of AI software within an AI-enabled legal services, therefore suggested to me that the attributes and performance of the technology would affect the level of trust legal professionals would report in an AI-enabled legal service.

The above distinction between human trust and trust in technology, is further complicated by research in the 'computers are social actors' paradigm, which indicates humans perceive different levels of 'humanness' between technologies (Li, Hess and Valacich, 2008), meaning a single set of trust beliefs in technology is not appropriate (Lankton, McKnight and Tripp, 2015). This is particularly relevant to trust in AI, where AI systems are sometimes designed to mimic human behaviour; meaning it is an open question whether the trustworthiness of AI is evaluated using beliefs relating to human trust or trust in technology. In the context of my own research, this meant there was an opportunity to see whether the legal professionals using AI in their work described their experience of interacting with AI using concepts from the human trust or trust in technology literatures.

The wider environmental context that both humans and technology are situated within, has also been found to mediate the development of trust (Gefen, Karahanna, and Straub, 2003;

Schaefer *et al.*, 2016). *Situation normality* refers to environments that are judged as being typical and in line with expectations. This increases the likelihood of trust developing as users are more certain of what the outcome will be of interacting with an information system. For example, a legal professional being asked to use AI would be more likely to trust the AI if the wider process it was situated within was familiar to them and reflected the wider norms of professional practice. A second aspect of context, *Structural assurances*, refers to the safeguards that surround an information system. In the context of an AI-enabled legal service this might be the knowledge that the firm's underwriters have agreed that use of the technology falls within the cover of the firm's liability insurance. Together *situation normality* and *structural assurances* are referred to as *institutional antecedents* of trust.

Environmental factors affecting trust also include those relating to the nature of the *task* that the system is performing. Where there is judged to be a high-level of complexity and/or risk involved, trust in the system is reduced (Schaefer *et al.*, 2016). The provision of legal advice can require the expertise of several different legal professionals, meaning the legal decision-making process can be complex and opaque; this is only increased further with the addition of technology such as AI. Furthermore, the often high stakes attached to legal matters, combined with significant ambiguity between legal advice and its outcomes, means the environmental context is inherently risky for those involved. Taken together the context of the legal sector means environmental factors are likely to play a significant role in determining the level of trust in AI, relative to other organisational contexts.

A further environmental factor affecting trust is *collaboration* (Hancock *et al.*, 2011). Information systems require their constituent parts to work together effectively to produce an output meaning collaboration between the different elements of the system is important. The absence of effective collaboration will, therefore, reduce the likelihood of the system being trusted to perform as expected. In the delivery of AI-enabled legal services, collaboration is required amongst legal professionals, potentially from different legal and technical disciplines; and between legal professionals and the AI software that underpins the process. Hence the level of coordination achieved in the delivery of an AI-

enabled legal service is likely to have a significant impact on the extent to which legal professionals develop trust in the system.

Researcher reflections

My review of the trust literature highlighted to me the complex and interdisciplinary nature of research in this area, and that trust was an important concept when seeking to understand why individuals choose to use a technology or not. The models of trust I reviewed displayed some commonalities, but it was not evident to me that any of the existing attempts to develop an integrated model of trust would provide a suitable a priori theoretical framework for my research. A specific concern being that there was a lack of evidence to suggest that the empirical findings underpinning the various models, would generalise and reflect interactions between users and AI systems, given the distinctive nature of AI as a technology (Butcher, 2019) or the specific context of the legal profession.

However, the relative consensus surrounding the role of the three broad factors – human, technological and environmental to the development of trust, led me to conclude that these were likely to play a role in explaining the development of trust in AI-enabled legal services. Furthermore, this also indicated to me that my approach to data generation would need to be flexible enough to allow each of these three factors to be considered. This suggested a more open-ended approach to data generation would be required, that did not focus exclusively on one of these factors at the expense of the other two.

2.5 Methodological implications of the initial literature review

It is also important to reflect on *how* best to study AI-enabled legal services. As a novel technology, it should not be assumed that traditional research methods necessarily offer the best means to study AI-enabled legal services. For example, Bailey and Barley (2020) advocate undertaking a more unified approach to AI research that goes beyond issues relating to design and use. Below I outline the relevance of the points they raise to studying AI-enabled legal services. Having done this I then discuss the methodological literature I consulted, in response to my reflections on how best to research AI-enabled legal services, which were outlined earlier in this chapter.

Bailey and Barley (2020) argue that studying technology at a high level of abstraction means the consequences that arise from the interaction of the technology and its context are overlooked. This critique reflects more general criticisms that IS research is over-reliant on experimental research using university students (Schlichter and Rose, 2013), and that cultural factors are frequently overlooked in user-system interaction research (Vance, Elie-Dit-Cosaque and Straub, 2008). Methodologically this highlights the importance of using field-based research methods, while at the same time recognising that it may not be valid to generalise research findings across contexts. In seeking to overcome these limitations, an increasing number of IS researchers are now adopting the use of case studies, underpinned by the paradigm of critical realism to better understand technological phenomena in their organisational context (Wynn and Williams, 2012). Having been alerted to this trend I familiarised myself with examples of how this research design has been used to conduct empirical research, noting its application to explain a range of complex issues, such as IT-led organisational change (Henfridsson and Bygstad, 2013; Volkoff and Strong, 2013). While the legal sector rarely features in IS research, its distinctive features (Empson *et al.*, 2015), increase the likelihood that the use of AI-enabled services in a law firm context will display distinct differences to its use in other organisational settings. This suggested to me that a case study approach could be well-suited to investigating AI-enabled legal services. Ultimately, this led me to focus on both a distinct phenomenon (AI-enabled legal services) and a distinct context (UK-based commercial law firms), using a case-based methodology capable of exploring the interaction between the two (outlined in Chapter 4).

Overlooking the impact of power is the second limitation Bailey and Barley (2020) highlight. They suggest stakeholder power is rarely scrutinised in research focusing on IS use, despite the fact that power dynamics will affect user acceptance in ways that are distinct from those factors that relate directly to the technology itself. In the context of AI-enabled legal services, there are several sources of power that might affect levels of AI use. Formal power is typically vested in the firm's Partnership who, as owners of the firm, are responsible for determining its goals and making investment decisions about the extent to which AI is used to further these goals. The power to design/procure and deploy AI-enabled legal services is typically delegated to subject-matter experts working in the multidisciplinary teams identified by Armour and Sako (2020). Finally, individual legal professionals can also exert

power if they are given the autonomy to decide whether to incorporate AI into their own professional practice. Methodologically this highlighted to me the importance of generating and analysing data using methods that were sensitive to issues of power (Kelemen and Rumens, 2008; Myers and Klein, 2011), and that my research participants needed to be drawn from a range of different stakeholder groups. The specific impact this had on my sampling strategy and methods choice is discussed in Chapter 4.

Alongside overlooking the role of power, IS research has also been guilty of ignoring the role of ideology (Bailey and Barley, 2020). Differences in stakeholder ideology can have a significant impact on technology, affecting both its design and use. The ideology of AI developers is shaped by their background in the computer and information sciences. Legal professionals in contrast have an ideology shaped by the organisational logics highlighted earlier in this chapter. In the context of AI-enabled legal services, these differences get reflected in the degree to which different stakeholder groups privilege the technical over the social when making choices about technology use (Forsythe, 2001). Tensions arising from these differences might manifest in relation to decisions such as: whether AI should be used as a substitute for legal professionals or merely something to augment their roles. Ideology may also impact the goals AI is programmed to work towards, for example is the goal of AI to increase human productivity or promote human well-being? Once again, the issue of ideology highlighted to me the importance of engaging with multiple stakeholders, using methods that could capture the subjective perceptions of individuals, for example their understanding as to the motives behind AI-enabled legal services being introduced in their firm. This led me to identify the use of more flexible, participant-led methods for data generation.

In addition to the above, Bailey and Barley (2020) identify the impact that AI can have beyond the boundaries of the organisation in which it is deployed. In relation to AI-enabled legal services this would mean thinking about what the wider impact of AI might be beyond the law firm and its employees. For example, how AI might impact the firm's clients and supply chain, or the legal system as a whole. Methodologically this highlights the value of research that extends beyond organisational boundaries; however, the scope of such research makes it challenging for individual researchers to conduct, and so I made a

pragmatic decision to not let this specific concern affect the methodological design of my research.

Interviewing objects

Beyond the points raised by Bailey and Barley (2020), are the methodological issues that stem from the challenges AI poses to the technology as a tool perspective. This means recognising the possibility that both humans and AI might be better understood as independent actors, whose interactions affect one another and lead to the creation of an assemblage. Hence, to understand complex socio-technical systems like AI-enabled legal services, there is a need for research methodologies that create opportunities to see both the human and the technology at once, rather than privileging one at the expense of the other.

To understand how to reflect this in my own research, I familiarised myself with recent developments in qualitative methods designed for a digital environment (Paulus and Lester, 2021). This made me aware of how my own use of technology as a researcher needed to be carefully considered, and informed how I might best generate data about AI-enabled legal services. This led me to identify *Interviewing objects* (Adams and Thompson, 2016), a research methodology designed to allow digital objects to be treated as research participants. The approach aims to give these objects 'a voice of their own', thereby makes them available for critical analysis. This is achieved through the use of different heuristic devices that enable researchers to pay close attention to the everyday interactions that digital objects have with both humans and non-human entities. The aim is to help researchers attend to both the material world and social world simultaneously.

Adams and Thompson (2016) offer eight different heuristics to *interview objects* but stress that they should not be used in a prescriptive fashion. Thinking about the goals of my own research I selected six heuristics, which I judged would be useful in helping me generate the data I required. The chosen heuristics are briefly described below (see Table 6), with an example of how the heuristic was applied to generate relevant data.

Table 6: Heuristic devices chosen to generate data

Heuristic	Description
Gathering Anecdotes	<p>Recreate the everyday events and situations in which humans and objects participate. They can be first-person experiential and third-person observational.</p> <p>Interview question: <i>Describe how AI-enabled legal services became part of your professional practice?</i></p>
Following the actors	<p>Identifies actors relevant to the phenomenon of interest, and studies the ways in which related micro-practices are performed.</p> <p>Interview question: <i>Which AI-enabled legal services do you have personal experience of? Tell me about some of the specific practices linked to the delivery of these services? Who or what is involved?</i></p>
Listening for the invitational quality of things	<p>Considers what the affordances of the technology invite users to do, think, or perceive; or conversely what they discourage.</p> <p>Observation: <i>What types of interaction does the design of the AI software's interface encourage or discourage?</i></p>
Studying breakdowns, accidents and anomalies	<p>Identify what only becomes apparent in moments when a technology fails to work as anticipated and, therefore, becomes visible to us.</p> <p>Interview: Are you aware of any instances when an AI-enabled legal service has failed to work as expected?</p>
Discerning the spectrum of human-technology world relations	<p>Determine what kind of relationship exists when humans and technology interact with one another.</p> <p>Interview: I'm going to describe four common ways in which people can relate to technology. I'd be interested to know if one or more of them describes how you relate to AI, or whether you would describe your relationship differently.</p>
Applying the 'laws of media'	<p>Reveal the ways in which humans understand a technology to affect its surroundings.</p> <p>Reflective exercise: Capture how AI-enabled legal services affect your work. What human capacities are enhanced or diminished?</p>

2.6 The opportunity to make a contribution

My initial literature review provided me with a number of useful concepts and theories of relevance to my research. Research conducted within the professional services field alerted me to the frameworks and models most commonly used to explain a range of organisational phenomena, within the specific context of the professions. While technology-related

phenomena have not historically been a major focus of the field, it was encouraging to see an increase in the number of papers relating to the impact of technology on the professions, with the growing use of AI an important element of this. This confirmed to me the relevance of my research topic to academics within the field. With most of the extant research focused on technology's impact at the organisational level I was also confident that I could make a contribution by focusing my research on the impact of AI at the level of individual professionals. This would enable me to add knowledge by more clearly distinguishing AI-enabled legal services from similar phenomena, and focusing on how AI was impacting the practices of legal professionals.

My review of the IS literature made me aware of important issues to consider when a new technology is introduced, and the theories that have been developed to study this, for example user acceptance and trust; concepts that were not readily apparent in conversations about technology within the professional services literature. It was also helpful to understand how IS researchers were beginning to question the ability of their existing theoretical models to explain AI as a phenomenon, and what features make it distinct from other technologies.

Having drawn on the IS literature to identify several ways in which AI-enabled legal services might not conform to traditional understandings of user-system interactions, I was confident that my research could make a number of contributions. First, I could generate data that would indicate whether the technology within an AI-enabled legal services was understood to be a tool by the legal professionals who used it. Second, I could explore how trust in AI did (or did not) develop amongst legal professionals. Third, my research would look to explain the reasons why legal professionals in UK law firms did (or did not) use AI-enabled legal services. Together these aims gave me three exploratory research questions that would guide my approach to generating data, but which would remain open to revision, depending on the data that was generated:

1. *How are AI-enabled legal services understood by the legal professionals that use them?*
2. *How do legal professionals develop trust in AI-enabled legal services?*
3. *What explains the use and non-use of AI-enabled legal services amongst legal professionals?'*

The process of revision that led to these initial research questions being refined and reduced to the two over-arching research questions that featured in the Introduction, is discussed in Chapter 5.4.

While the results of my research would not be suitable to generalise across contexts, this type of localised theoretical explanation would be valuable to both researchers interested in the effect of AI on workers in professional service firms, and organisations wanting to promote the use of AI amongst their workforces.

2.7 Revisiting the literature during my research

During the process of data generation and analysis, I sought further theoretical insights to help me explain my findings. This was particularly relevant when the patterns identified in the data deviated from what my initial review of the literature had led me to anticipate might be found. This prompted me to consider further concepts and theories which I did not identify during my initial literature review. Their relevance to the study of AI-enabled legal services is now discussed. In addition to this, I also sought to keep abreast of newly published research relating to the use of AI within different organisational contexts, including professional service firms. In doing so, I noted that the increased use of AI had been accompanied by growth in the extant knowledge of AI across the research streams discussed in the initial literature review.

Revisiting the literature helped me identify ongoing challenges within AI research that my own research had an opportunity to help address. I also became more aware of the relevance of research focusing on how knowledge work is impacted by technological change. While the impact of technology on knowledge work has been studied for some time, there has been an increase in research which considers the specific impact of AI on knowledge work. This is the research stream that I determined the findings of my own research could most significantly contribute towards, I therefore, use this further review of the literature to situate my research within this broader discussion.

Ongoing challenges in AI research

The aforementioned challenges AI poses to research framed by the technology as a tool perspective (Schuetz and Venkatesh, 2020) have continued to influence AI research.

Ågerfalk *et al.*, writing in 2022, call on researchers to reconsider the traditional conceptualisation of *agency* as a uniquely human quality, suggesting that AI systems that use machine learning can be better understood as digital agents that are capable of enacting and shaping institutions and their institutional logics, rather than passive tools. This means AI should not simply replace more general conceptualisations of IT in existing research models; and furthermore it is important to distinguish between different types of AI, rather than treat it as a homogenous technology.

Anthony *et al.* (2023) agree with this assessment arguing that the usefulness of extant technology-focused management research remains limited because of the ways in which AI differs from earlier technologies. Specifically, they highlight AI's lack of functional consistency and functional transparency, and its invisibility to users (three of the five characteristics previously identified by Schuetz and Venkatesh, 2020). They conclude that an alternative research perspective is required, which treats AI as an actor within a system, and which focuses on how AI and human actors interact with one another.

While there is increasing evidence of alternatives to the tool perspectives being used to investigate the use of AI in the context of the UK legal sector – for example, Spring, Faulconbridge and Sarwar (2022) adopt an 'ensemble' view of the relationship between technology and professional – such perspectives are not widespread (Glaser, Sloan and Gehman, 2024). The methodological challenges Anthony *et al.* (2023) highlight to researchers wishing to adopt their recommended 'AI as a counterpart' approach, are largely addressed by Adams and Thompson's (2016) approach to 'interviewing objects'. I concluded, therefore, that the challenges identified in my initial literature review remained relevant, and by focusing on the interactions between AI and legal professionals, my research was well-placed to contribute to this ongoing debate. This led me to identify and utilise two further theories during my analysis – technological frame theory and material hermeneutics – that did not feature prominently in my initial literature review, but which had the potential to offer valuable insights into the nature of the relationship between legal professionals and AI.

Technological Frame Theory

Socio-cognitive approaches to human-computer interactions, posit that human behaviour towards technology can only be understood through reference to how that technology is first interpreted. The concept of the technological frame seeks to capture the outcome of this interpretive process, with the frame reflecting the knowledge, assumptions, and expectations that an individual or group has about a particular technology (Orlikowski and Gash, 1994). A technological frame can, therefore, be understood as a simplified mental model of the technology, which by directing an individual's attention in order to assign meaning, has the potential to influence individual behaviour relating to technology development, implementation and use (Elbanna and Linderoth, 2015). This makes technological frames of relevance to organisations seeking to develop strategies that will facilitate the introduction of new technologies, such as AI (Spieth *et al.*, 2021).

Technological frames reflect not just the material dimensions of a technology but also what it symbolises to those who interact with it. In those instances where technology carries heavy symbolic value, it is particularly important that the claims made about a technology are considered alongside its technological features in order to explain the ways in which it is used (Anthony *et al.*, 2023). This means technological frames are also shaped by the individual's personal experiences and affiliations (Cornelissen and Werner, 2014); and the specific organisational and institutional context, within which the technology is experienced (Davidson, 2002).

Based on a review of the literature relating to technology and organisational change, Orlikowski and Gash (1991) initially identified seven potential dimensions of a technological frame. However, they cautioned that the list was unlikely to be exhaustive, and that the structure and content of a technological frame would be dependent upon its local context, meaning frames should not be expected to generalise. This was reflected in their subsequent research, which identified in situ technological frames comprising of only three different dimensions (Orlikowski and Gash, 1994). Hence, they argue the structure and content of a technological frame needs to be examined in situ rather than determined a priori.

As the literature of technological frames has grown, the opportunity has arisen to investigate whether there is empirical evidence suggesting common frame dimensions across contexts. An early review (Davidson and Pai, 2004) found that while the precise names given to frame dimensions did tend to be context-specific, there was evidence of similar dimensions emerging across settings, in addition to novel domains that were context-specific. For example, several of the studies they reviewed identified frame dimensions similar to what Orlikowski and Gash termed *the nature of technology*, which related to an individual’s understanding of a technology, its capabilities and its functionality. More recently, Spieth *et al.* (2021) identified five dimensions of technological frames, following a factor analysis based on a review of over 19 published papers (see Table 7).

Table 7: Dimensions of technological frames

Dimension	Description
Personal attitude	Whether the individual has a positive or negative attitude towards the technology; their level of expectations; and the extent to which technology plays an important part in their life.
Application value	The extent to which technology has a positive or negative impact on important aspects of working life, e.g. working flexibly, producing work that is accurate and high quality.
Organisational influence	The extent to which an individual’s colleagues/organisation is supportive of technology use and makes its use easier (or mandatory).
Industrial influence	The degree to which technology is used within an industry; and whether technology use is demanded by suppliers, customers and competitors.
Supervisor influence	Whether the individual’s supervisor recognises the benefits of technology, is proficient in using it; and makes it possible to use technology to undertake work.

Source: Adapted from Spieth *et al.*, 2021

These five dimensions capture both those that can be seen to be rooted in the individual actor (*personal attitude* and *application value*) and those rooted in their affiliations both within and outside their organisational context (*supervisor influence*, *organisational*

influence and *industrial influence*). While the names given to the five emerging dimensions differ from those used previously, the framework provides a useful reference point for researchers when analysing their data, to identify similarities and differences in the technological frames of different individuals and groups.

In cases where individuals use multiple technologies there is scope for different frames to develop in relation to each technology, with the structure and content of a frame influencing an individual's behaviour towards the associated technology. Congruence amongst frames occurs when the various frames an individual holds broadly align in terms of their structure and content, although they need not be identical (Orlikowski and Gash, 1994). In such situations, congruence means an individual is likely to display similar behaviour towards different technologies within the same organisational context (Minkkinen, Zimmer and Mäntymäki, 2023). Conversely frame incongruence arises when technological frames exhibit a difference in kind (rather than degree), leading to variations in behaviour towards different technologies. For example, in the context of a law firm, the technological frame a legal professional develops about AI-enabled legal services, could differ significantly from the frame they hold about other technologies that they use, such as email. This might be explained through identifying differences in the *application value* of the two technologies; if the value attached to email was positive but the value attached to AI was negative then despite the shared organisational context, the legal professional may elect to use email but not AI.

Moving from the individual to the group-level, frame incongruence can arise when there are variations in frame structure or content amongst group members. For example, in the context of a law firm, the technological frame a qualified lawyer develops about AI-enabled legal services, may differ significantly from other professionals e.g. an IT developer responsible for the implementation of the AI-enabled legal service. This might be explained through identifying differences in their *personal attitudes* to AI, arising from variations in their professional training and the extent to which AI play a role in their sense of professional life; despite both individuals sharing in common the dimensions of *organisational influence* and *industrial influence*.

In contexts where the structure and content of a frame remains unchanged, technological frames can become taken-for-granted by groups and individuals, meaning they are rarely reflected upon, but their effect on behaviour persists. Changes that affect the structure or content of a frame, for example changing market conditions or the introduction of a new technology to a workplace, can trigger frame re-evaluation, leading to shifts in meaning, and ultimately behaviour, towards a technology (Davidson, 2002).

Technological frames are, therefore, a useful concept for researchers seeking to explain how different individuals and groups interpret and interact with technology, including AI although most research has tended to focus on public perceptions of AI rather than employees within organisations (Minkkinen, Zimmer and Mäntymäki, 2023; Wang and Liang, 2024). In the context of my research, there was, therefore, an opportunity to see whether the data suggested the existence of a common technological frame amongst legal professionals; or if there was evidence of frame incongruence amongst different groups. With any differences in the structure and content of their technological frames indicating which dimensions were shaping their decision to use of AI-enabled legal services.

Machine Hermeneutics

My process of data generation used heuristic devices suggested by Adams and Thompson (2016) to explore how legal professionals understood their relationship with AI, within the context of an AI-enabled legal service. Abbott's (1988) in identifying three types of professional act, identifies inference making as the act that is of central importance to professionals. This led me to anticipate that legal professionals might be most likely to identify their relationship with AI to be hermeneutic in nature (Ihde, 1990); as this would preserve the status of the legal professional, as responsible for generating meaning, and thereby creating valuable knowledge for clients. I, therefore, consulted literature focusing on hermeneutic relationships with technology.

Ihde (1990) argues hermeneutics needs to make the material world, rather than just text, an object of analysis. This allows recognition that the human act of interpretation is often mediated by the use of technology, such that it is not just human eyes that 'read', for

example, a scientist interprets a sample of human tissue by observing it through the lens of a microscope. In this way technology can contribute to the human act of interpretation.

More recently Ihde's ideas have been further developed with specific reference to AI (Hongladarom, 2020; Wellner, 2020). The former of whom argues that we have now entered a period of 'machine hermeneutics', in which machines are now capable of undertaking the hermeneutic act independently of humans. This is due to the ability of AI algorithms to read and process data; and attribute their own meaning to the data, before acting upon this interpretation. This means an interpretive act has taken place prior to the processed data being presented to a human, who through their own reading of it can then add a further level of interpretation. This is significant as while it can be argued that the AI algorithm can help humans to understand reality better e.g. by presenting the data in a more understandable format; it is also plausible that algorithmic interpretation can distance humans from reality and cause humans to be misled. The relevance of this to the work of legal professionals has recently started to be acknowledged, with a recognition that the output of AI software underpinned by machine learning is often opaque to the professionals who act on the output, due to their inability to see the inferences being made by the software (Faulconbridge, Sarwar and Spring, 2024). This more complex understanding of the hermeneutic relationship was conceptually useful to me as it offered the opportunity to more fully explore the nature of the relationship experienced by legal professionals when using AI when analysing my data. This would provide further insights into the process through which legal professionals decided to use AI-enabled legal services.

Technology and the changing nature of professional work

In seeking to better understand the potential mechanisms through which AI had the potential to impact the practices of legal professionals, I consulted the literature relating to the impact of AI (and technology more generally) on knowledge work and the professions; an area of research that has developed rapidly during the period of my own research.

The term knowledge work refers to the application of specialised theoretical and analytical knowledge to non-routine problem-solving and the development of related products and services. Such work is typically undertaken by highly educated individuals, who using their domain-specific skills and expertise are able to generate work that is of significant value to

the clients (Alvesson, 2001). Knowledge work has long involved the use of specialised tools and instruments, through which knowledge workers are able to demonstrate their expertise. This means the effects of technology are typically mediated by the professionals using them, they do not happen automatically (Kellogg, 2022).

In recent decades the practices of knowledge workers have become increasingly dependent on the use of digital technology with AI now beginning to impact the practices of knowledge workers in significant ways. While traditional *production technologies*, such as word processors, have increased the efficiency of knowledge workers, they have done so without fundamentally changing the ways in which knowledge workers generate knowledge. In contrast to this, AI is an example of what can be more accurately understood as an *epistemic technology*, whose primary purpose is to enable the generation of knowledge (Anthony *et al.*, 2018). Such technologies significantly enhance the analytical capabilities of individual workers, meaning they can have a transformational effect on professional practices, for example the use of computer-aided design software by architects to design buildings. However, the use of epistemic technologies can also distance workers from key tasks that they previously would have undertaken using more manual tools and techniques – this presents a risk to both professional expertise and the output of knowledge work processes (Anthony, 2021).

During previous waves of technological change, knowledge workers with significant levels of expertise have responded to the introduction of epistemic technologies by seeking to develop a deep understanding of the functionality of the technology and how it works (Anthony, 2018). This was achieved through engaging in practices such as peer review and comparing the output produced by subject-matter experts using traditional methods, to the output produced using new technologies. This understanding allowed knowledge workers to develop trust in the technology; ensure the ongoing accuracy of their work; and preserve their expert status (Bailey and Barley, 2011). In contrast, it was suggested that non-experts who were more socially distant from the technology, such as general managers, would be less concerned about understanding the inner workings of the technology used in their organisation (Bailey, Leonardi and Barley, 2012). In summary, users with expert knowledge would *question* their technology, while those with more limited expertise would *accept*

them (Anthony, 2018). These distinctive patterns of behaviour indicate the importance of both a worker's expertise and their structural position relative to the technology.

The black boxing of technology

As previously discussed, the complexity of the algorithms that underpin AI makes it increasingly difficult for the professionals who interact with AI to understand the technology and, therefore, critically assess the knowledge that AI produces (Anthony, Bechky and Fayard, 2023). This can result in AI essentially being treated as a 'black box' that lacks transparency. Under such conditions it might be anticipated that knowledge workers who find themselves unable to question and develop their understanding of AI, will be unwilling to defer to its decisions, especially when they contradict their own professional judgement. Alternatively, experts may feel they have no choice but to trust AI, with their trust reflecting their trust in the designers of the technology, rather than their own understanding of how it works (Anthony, Bechky and Fayard, 2023).

Empirical research within the investment banking sector by Anthony (2021) has demonstrated that, contrary to earlier research, the level of technological understanding amongst knowledge workers (in this case bankers using AI software for financial analysis) can vary significantly, even amongst close colleagues who use identical technologies to perform the same job role. The findings of this research indicate that while some knowledge workers do continue to engage in questioning practices that seek to unpack the assumptions that underpin AI, others adopt accepting practices, which it was previously assumed were characteristic of non-experts. These accepting practices allow the output of AI to be treated as a legitimate part of the process of knowledge production, but without the need to understand how the technology works. Instead these professionals sought to validate the knowledge produced by AI-enabled processes by scrutinising the output of the process, rather than the technology itself. Anthony (2021) explains the observed behavioural differences through reference to the different ways in which the analytical tasks performed by the bankers were split across different team members. The findings of this research, therefore, suggest that contextual factors relating to task allocation and hierarchy can have a significant impact on the ways in which professionals seek to validate the output of AI and how they develop trust in it. Hence, while expertise may be necessary in order for

professionals to developing a deep understanding of complex technologies such as AI, expertise alone is not sufficient to ensure that this happens in all contexts.

Reflecting on these findings, it occurred to me that the different kinds of validating behaviours could be seen to map onto two of the three factors identified in my initial review as being linked to the development of trust in AI. Questioning behaviours can be seen to develop trust through understanding the technology-related factor, while accepting behaviours lead to trust development through the environmental factor. While both approaches to validation can lead to professionals trusting and using AI, the process through which trust develops is different in each case. This is non-trivial, given the high levels of trust that wider society places in the judgement of professionals, who it is assumed possess expert knowledge of their field and its associated practices, including the use of relevant technologies.

In the context of my own research, legal professionals are constantly generating knowledge through their work – knowledge that carries great significance to their clients. In seeking to understand how the legal professionals in my research came to use AI in their work, the concept of validating practices (questioning and accepting) had significant potential to explain how trust might develop amongst different groups of legal professionals. I was also aware that the specific contextual conditions and mechanisms identified in Anthony's (2021) research could not be assumed to generalise to the legal sector, meaning that there was an opportunity to develop a novel theoretical explanation of AI use amongst lawyers, based on their engagement (or not) in the practices of questioning and accepting.

To help me develop a causal explanation of the interactions of legal professionals with AI, I consulted the wider literature to see what other factors had been found to influence how professionals interact with technology through their professional practice. These insights would be useful in identifying potential mechanisms that might be evident in my own data.

The coordination and protection of professional work

As discussed in the initial literature review professionals seek to define and organise themselves at the macro level, using formalised methods to establish jurisdictional boundaries around their expertise and their work, in order to maintain their privileged

status. However, this protective behaviour has been seen to be challenged by financial-economic, socio-economic and technological pressures (Noordegraaf, 2020). The hybrid model of professionalism proposed by Noordegraaf (2015) has been recognised as challenging to implement in practice, as it left organisations managing tensions between the competing logics of professionalism and managerialism. This has meant it has not always been able to address the aforementioned external pressures facing professional organisations. This has led to further models of professionalism being suggested.

Connected professionalism (Noordegraaf, 2020) goes beyond their original idea of hybridity, with a view to explaining how perceived professional boundaries can be overcome, so that professionals can work together in ways that help them meet the wider demands placed upon them. To achieve this, the model adopts a relational lens (Anteby, Chan and DiBenigno, 2016), emphasising the importance of openness and relationships, in contrast to earlier models that gave primacy to autonomy and closure. This requires professionals to recognise one another as being equally knowledgeable, independent, and authoritative within each profession's own area of expertise. Relationships between professional groups are, therefore, maintained in a system of distributed expertise and authority, reflecting the interdependent nature of the work they undertake together. These interdependent relationships are also recognised to extend beyond organisational boundaries, to include other important stakeholders, such as clients (Noordegraaf, 2020). In response to this Faulconbridge, Folke Henriksen and Seabrooke (2021) have critiqued the model of connected professionalism, by considering how well the model can explain recent developments within professional service firms using AI. While they identify changes in professional practice that are supportive of there being increased levels of connectedness, for example an increase in the importance of relationships between professionals and clients, they also note that these same connective practices can be understood as a strategic response, designed to protect professionals, what Faulconbridge terms *protective connectedness* (Faulconbridge, Folke Henriksen and Seabrooke, 2021). In the context of my own research both these concepts (connected professionalism and protective connectedness) provided useful frameworks through which to consider the behaviour of legal professionals using AI-enabled legal services.

At the micro level, jurisdictional boundaries can appear less formalised, meaning in situ negotiations between professionals can configure individual roles, determining who performs what tasks, and how the actions of different professionals are coordinated (Abbott, 1998). The emergent outcome of such localised processes will reflect the interplay of formal logics and situated practices, within the local context (Pine and Mazmanian, 2017), with accountability, predictability and common understanding important conditions that determine the nature of the coordinating that emerges (Okhuysen and Bechky, 2009).

The introduction of new technology-backed work processes, such as AI-enabled legal services, can be understood as an attempt by organisations to control such emergent outcomes, through the imposition of 'best practice' which has been determined in advance, and reflected in the design of the system. These new work processes can lead to role reconfiguration and more formalised methods for coordinating the work of different professionals e.g. interactions between professionals may be digitally-mediated and recorded, rather than being self-managed by the professionals involved. Such change can, therefore, play a role in how the jurisdictions of existing professions change over time, and the emergence of new professions .

The impact of technological change on the coordination of professionals has been found to vary – in some instances coordination has been negatively impacted, while in others coordination has been enhanced. For example, Pine and Mazmanian (2017) highlight the negative impact that an electronic health records system had on the ability of health professionals to coordinate their work. Kellogg (2022) highlight the role of power in determining such outcomes, suggesting that in many instances powerful professionals can respond to technological change by reconfiguring the work of less powerful colleagues in ways that protect their own role and status, but which may come at a cost to others. However, Kellogg (2022) does acknowledge that technological change can be mutually beneficial to different groups in contexts where those involved are given the opportunity to experiment with the technology and adapt it to their local needs i.e. the process is emergent rather than pre-determined. This suggested to me that it would be important to generate data that would help explain the extent to which AI-enabled legal services prescribed the working practices of legal professionals.

Task allocation in AI-enabled legal services

Anthony (2021) highlights how task allocation influences the type of validating behaviours that experts working in the banking sector engage in, in order to determine whether to use AI in their working practices. I, therefore, felt it was plausible that task allocation within AI-enabled legal services could help explain the behaviour of legal professionals in my own research. Research on task allocation within the legal sector has been published during the course of my studies, with the concept of *consumers* and *producers* of AI-enabled legal services proposed by Armour, Parnham and Sako (2022) in their research on how AI is used to augment lawyering. They highlight the importance of multi-disciplinary teams in bringing together the necessary human capital required to create the 'technology pipelines' capable of developing and managing AI-enabled legal services (Armour and Sako, 2020).

They identify two ways in which legal professionals can interact with AI technology within the context of a technology pipeline: as consumers or producers. Qualified legal professionals in fee-earning roles, whose work is augmented by AI, are described as *consumers*. In contrast, professionals working in multi-disciplinary teams and who possess legal knowledge (but who are not necessarily qualified to practice law) are identified as *producers*. The precise nature of the tasks that constitute the work of producers will vary according to their individual expertise but is likely to involve activity in all five stages of the technology pipeline (Table 8).

I first encountered the concept of the technology pipeline and the roles of consumers and producers during the period of in which I was generating data for my research. The technology pipeline model reflects an idealised approach to the development of AI-enabled legal services and is intended to describe the approach taken by large commercial law firms, such as those that featured in this research. This meant my research provided an early opportunity to generate empirical data that the model could be compared against and if necessary refined.

Table 8: Steps in a ‘technology pipeline’

Step	Description
1. Requirements	Collection of end-user requirements for the system that can be translated into a set of technical system specifications.
2. Design and procurement	Specifications are mapped to existing Legaltech offerings via a procurement process to identify a suitable system; or are used by IS developers (external or in-house) to develop a new system.
3. Data ingestion	After deployment relevant data is uploaded to the system for training purposes.
4. Data labelling	Ingested data is labelled by subject matter experts; feedback is provided to the system based on the output it produces.
5. Application of results	The trained model is applied to the use-case identified in the requirements; system output needs to be evaluated for significance, its implications determined and explained to relevant parties.

Source: Adapted from Armour, Parnham and Sako, (2022)

On reading the descriptions of the consumer and producer roles I was immediately able to recognise examples of individuals performing these roles and their associated tasks from my own professional experience of working in the legal sector, prior to commencing my research. However, my professional experience also suggested to me that the clear-cut distinction the model made between the two roles was an over-simplification, as I was personally aware of individuals whose job role comprised aspects of both producers and consumers. In addition, while Armour, Parnham and Sako (2022) provide detailed descriptions of *what* consumers and producers do, their research did not explain *who* they are or *what* their understanding of AI-enabled legal services might be. I, therefore, felt my research offered an opportunity to explore the concept of producers and consumers further, and better distinguish the different ways in which legal professionals in these roles relate to AI-enabled legal services. More specifically, I wanted to understand whether there was evidence of differences in the validating behaviours of producers and consumers, which might help provide an insight as to the mechanisms that could explain how legal professionals come to use AI-enabled legal services.

Following from the above, I wanted my research to be able to account for individuals whose work was comprised of tasks drawn from both the producer and consumer roles. The existence of professionals whose work spans the producer-consumer boundary is not anticipated by Armour, Parnham and Sako (2022), and so their existence in the data would require the model to be further developed to reflect this. There was also an opportunity to understand whether individuals in such roles were subject to any competing demands that might exist between the two roles, and if so, how these tensions were managed by them. The concept of a liminal space, which exists at the boundary of existing social structures (and yet is not fully part of either) or where new structures are emerging seemed a useful concept for understanding the experience of these professionals (Dale and Burrell, 2007). Individuals who inhabit liminal spaces typically find the experience ambiguous, as there can be an absence of the norms and practices found within more dominant spaces. While this can be unsettling and disruptive, it can also be a space of creative possibility, during which existing professional identities can be temporarily set aside and new ones explored (Chen and Reay, 2021). In navigating liminal spaces, individuals can, therefore, be understood to experience a three-stage 'rite of passage' (Turner and Abrahams, 2017), which sees them initially separate from their previous role; enter a state of social limbo; before finally experiencing incorporation, which sees them assume a new and recognised role or status. Liminality has been used to explore the experiences of individuals whose roles exist at the boundary of one or more professions (Paton and Hodgson, 2016; Empson, 2017), and explain how professionals can positively respond to changes in their working practices (Chen and Reay, 2021). It therefore struck me as a useful concept that could be used to describe the position of any participants that could not be straight-forwardly identified as producers and consumers. The process of liminality also seemed to have the potential to help explain the experience of such individuals, depending upon where they found themselves on their journey of liminality.

The role of clients in knowledge work

The impact of clients on how professionals understand and conduct their work has recently begun to attract increased levels of interest. While the professions literature has previously assumed that professional work is primarily shaped by forces from within the profession

(Abbott, 1988), more recent research indicates that professionals also adapt their working practices by responding to external forces, and in particular clients (Bourmault and Anteby, 2023). This was apparent in my own research where the data indicated that clients were understood to be a key driver of the adoption of AI-enabled legal services.

Bourmault and Anteby (2023) suggest clients can impact the behaviour of professionals through two different pathways. The first pathway, 'turning inward', reflects a professional's personal experience of how a client responds to a new practice. For example, a legal professional might observe whether a client responds positively or negatively to being told the firm intends to use AI to perform a due diligence exercise on their behalf. If the client responds positively, indicating the new approach is preferable to historic practices, the legal professional may find it easier to accept AI as part of their professional practice. In such instances, client behaviour can be seen to influence how professionals both understand and perform their work. The second pathway, 'turning outward' focuses on the nature of the relationship a professional has with their client. For example, a law firm may provide clients with a technology platform that allows them to 'self service' certain types of legal work e.g. a client is able to select their preferred form of contract from a suite of precedents, rather than the legal professional choosing on their behalf. Changes of this sort might cause a legal professional to reevaluate their understanding of what type of working relationship a client expects from a legal professional. In seeking to understand why legal professionals chose to accept the use of AI, I found the concept of 'turning inward and outward' helpful in explaining why legal professionals might understand external factors to be driving the use of AI within the profession.

2.8 Conclusion to the Literature Review

The above review of the literature demonstrates how researchers have explored the use of AI within the workplace in recent years. The review highlights the different ways in which the use of workplace AI has been researched and theorised across different domains; and in doing so, identified useful concepts and theories to inform my own research.

More broadly, the extant research indicates that technology (but not specifically AI) is frequently used by professionals to generate knowledge and demonstrate their expertise, provided the technology is understood and can be trusted (Bailey & Barley, 2011). Such

generalised findings have, however, been challenged by more recent AI-focused research (Anthony, 2021), which suggests professionals will use workplace AI despite not being able to access and understand its inner workings. These findings highlight the importance of localised contextual factors in explaining AI acceptance.

Research in the Information Systems field has also taken significant interest in micro-level user-system interactions, with theoretical models such as the Technological Acceptance Model (Davis, 1989), seeking to explain how individuals come to use different technologies. However, it has been recognised that the cognitive focus of such models can neglect the role of wider contextual factors in shaping technology use (Jensen and Aanestad, 2007). Furthermore, the theoretical assumptions that underpin several IS models of trust in technology and models of technology acceptance, are challenged by AI (Schuetz & Venkatesh, 2020; Anthony *et al.*, 2023). Hence, it is suggested that AI researchers may need to develop their own models, rather than simply interchange AI with the more generalised conceptualisations of the IT artifact that appear in extant research (Ågerfalk *et al.*, 2022).

My research sought to address the aforementioned limitations identified within the extant research, by both exploring how AI is understood amongst legal professionals (rather than professionals in general), and generating a theoretical model capable of explaining AI use amongst legal professionals within their local context. The process of analysis that was used to answer these research questions was grounded in the empirical data, which was generated with professionals who possessed first-hand experience of using AI within UK law firms. The following chapters provide an overview of the research approach and strategy that was adopted to achieve these goals.

3. RESEARCH APPROACH AND STRATEGY

This chapter begins with an overview of the key tenets of a phenomenon-led research strategy and highlights both the relevance and implications of adopting such a strategy to researching AI-enabled legal services. I then explain how a critical realist theoretical standpoint was chosen to underpin the research, by demonstrating how it was better able to achieve the goals of the research than the alternative paradigms of positivism and interpretivism.

Having demonstrated consistency between the research strategy and theoretical standpoint, I then explain the broad methodological principles that underpin critical realist research, contextualising them through reference to researching AI-enabled legal services. I also describe how the chosen research design, a critical realist case study, satisfied these principles.

3.1 An introduction to phenomenon-based research

The domain of management generally requires that research makes a theoretical contribution, but there is a risk that such demands can deter researchers from investigating and describing important contemporary phenomena, about which theory is yet to be developed (Hambrick, 2007). As I have outlined previously, the nascent use of AI by legal professionals means research about the phenomenon, while growing, remains limited. I therefore chose to adopt a phenomenon-based approach in my research, the implications of which are described below, both in terms of the impact on my overall research strategy and selected methodology.

While phenomena are the subject of all research, phenomenon-based research prioritises the rigorous identification, description, and conceptualisation of a phenomenon, as a necessary precursor to more traditional theory-driven research. The emerging phenomenon is, therefore, regarded as occupying a middle ground between the empirical data and abstract theory. A phenomenon-based research process, therefore, looks to data to provide evidence of a phenomenon; from which theory capable of accounting for the

phenomenon can be developed (von Krogh, Rossi-Lamastra and Haefliger, 2012). Hence in phenomenon-based research, it is the phenomenon, rather than theory, which shapes the research process.

Management domain phenomena are typically complex, with outcomes reflecting the interplay of both social causes and environmental causes (that are independent of human behaviour). In such open systems, where individual elements can influence the whole system, the use of traditional experimental research methods (that seek to isolate specific variables within a controlled environment) is not desirable. Instead, the use of more open-ended research questions that can be answered through the identification of new concepts and variables is more appropriate (Daft and Lewin, 1990). This in turn means phenomenon-based research needs to adopt methods that can consider data generated at multiple levels – individual, group, organisation and societal (von Krogh, Rossi-Lamastra and Haefliger, 2012).

The phenomenon-based approach has a long tradition in management research and has contributed significantly to the development of knowledge and theory in many fields, for example, Mintzberg's (1978) research in the field of strategy. More recently, in the field of professional service firm research, phenomenon-based research has featured prominently in academic journals, for example, Pareliussen, Æsøy and Giskeødegård's (2022) research on the impact of technology in the maritime industry. These examples demonstrate the potential contribution that phenomenon-based research can make when it is undertaken rigorously and with a clear purpose.

3.2 A phenomenon-based research strategy for AI-enabled legal services

The likelihood of phenomenon-based research making a significant contribution is maximised when the researcher has a clear understanding of what such research can and cannot hope to achieve. Von Krogh, Rossi-Lamastra and Haefliger (2012) have developed a useful framework to help researchers reflect on the maturity of the phenomenon they are researching (and the likely opportunities where a contribution might be made) prior to selecting their research strategy. In my research I utilised this framework to reflect on the

available opportunities for researching AI-enabled legal services, and to identify a suitable research strategy. Below I outline the key elements of the framework and show how I used it to generate a research strategy to investigate AI-enabled legal services.

Three phases of study

The von Krogh, Rossi-Lamastra and Haefliger (2012) framework identifies three phases of study that management phenomena typically go through – *Embryonic*, *Growth* and *Maturity* – as interest in them develops amongst researchers. In the *embryonic* phase researchers focus on distinguishing a new phenomenon from other better-understood phenomena. Initially there may be disagreement amongst researchers as to whether the phenomenon represents something novel, or whether it is an example of an existing concept. The number of researchers actively researching the phenomenon is likely to be low, with their efforts typically uncoordinated and potentially spread across multiple fields.

The *growth* phase is characterised by increased interest amongst researchers as the phenomenon becomes more visible and better understood. Researchers often start to coalesce into groups with increased coordination and collaboration of research taking place. Any discrepancies that exist between the phenomenon and the extant theory seeking to explain it, serve as a stimulus for further research using more novel research designs.

Entrance to the *mature* phase reflects growing agreement about the nature of the phenomenon. Building on the earlier phases of research, researchers become better able to predict outcomes associated with the phenomenon, and consistency of findings is often reported across different contexts.

AI-related phenomena (including AI-enabled services) have recently attracted increased interest amongst academics working across a variety of fields, as demonstrated by recent systematic literature reviews on the topic (Çelebi, 2021; Collins *et al.*, 2021). While much of the initial research on AI took place within the domain of Information Systems, journal articles are now appearing with increased regularity in the domain of Management and the field of professional services. There is also evidence of researchers starting to coordinate their efforts and collaborate with one another both in terms of research projects, for example, the £20m UKRI sponsored Next Generation Services Challenge which ran from

2018-2022 (UKRI, no date) and conference streams, for example, *Professions in an age of intelligent technologies* at the European Group for Organizational Studies (EGOS, no date). While this indicates the phenomenon is starting to move beyond its embryonic phase, the volume and type of research taking place suggests the phenomenon is still some way from reaching maturity (von Krogh, 2018). This means analysis of AI-enabled services frequently takes place alongside other technological phenomena such as automation, block chain and virtual reality, as part of wider discussions about the use of technology within organisations.

Based on the above, I designed my research into AI-enabled legal services so that it was aligned to a phenomenon in its growth phase of study. In contrast to other recent research, which focused on broader topics, such as the digitalisation of professional service firms (Kronblad, 2020), I chose to focus my research on a single phenomenon (AI-enabled legal services), within a specific context (UK commercial law firms). This narrower approach meant I was able to consider the phenomenon in greater depth. While phenomena-based research on AI-enabled legal services has begun to emerge since I started my PhD studies, most of the research has considered the phenomenon at the macro level, for example, Armour, Parnham and Sako (2022). My research complements these papers, by instead focusing on how the phenomenon has been experienced and understood by the legal professionals who have made it part of their professional practice.

Five research strategies

Von Krogh, Rossi-Lamastra and Haefliger (2012) identify five different research strategies (*Distinguish, Explore, Design, Theorise* and *Synthesise*) that can be used either in isolation, or in combination with one another, depending upon the current phase of study and the goals of the research. A summary of the five strategies and their associated goals, is provided in Table 9.

To summarise, while a phenomenon is considered in the embryonic stage, exploratory work (*distinguishing, exploring* and *designing*) will characterise the work of researchers. Upon entering the growth phase, exploratory work continues but opportunities to *theorise* alongside this will increase. Finally, once the phenomenon is considered mature and a

significant body of research is available, more conventional research strategies that focus upon *theorising* and *synthesis* become the norm.

Table 9: Five phenomenon-based research strategies

Research strategy	Objectives of the strategy
Distinguish	<ul style="list-style-type: none"> • To give the phenomenon an identity by emphasising its distinctive characteristics, against the backdrop of existing practices • Establish conceptual boundaries that make it possible to consistently identify instances of the phenomenon.
Explore	<ul style="list-style-type: none"> • Collection of primary and secondary data relating to the now distinguished phenomenon. • Design of measures and frameworks to further sharpen the definition of the phenomenon and its boundaries.
Design	<ul style="list-style-type: none"> • Evaluation of existing research designs used to investigate the phenomenon to date. • Identification of new methodological approaches that can generate further insights and triangulate earlier findings.
Theorise	<ul style="list-style-type: none"> • <u>Embryonic/Growth phases</u> – focus on explaining puzzles that emerge in the empirical data through either refining extant local theory (developed within the specific context of the new phenomenon), or use of induction to generate theory from the available empirical data. • <u>Growth/Mature phases</u> – focus on identifying the extent to which local theories overlaps with (or deviates from) wider extant theory, to highlight similarities and differences with related phenomena. Alternatively, the focus can be to adapt, modify or combine existing theories so that they can better account for the new phenomenon.
Synthesise	<ul style="list-style-type: none"> • Minimise the needless repetition of research and overcome the fragmentation of knowledge and insights relating to the phenomenon across different fields. • Identify future opportunities for research about the phenomenon.

Source: Adapted from von Krogh, Rossi-Lamastra and Haefliger (2012)

Application of the framework to researching AI-enabled legal services

Having identified AI-enabled legal services as being in the early growth phase of research, I developed an overall strategy that encompassed three different activities. The first was to better *distinguish* AI-enabled legal services from other changes in legal practices, such as the more general use of technology by legal professionals. This would make it easier to determine whether established theories, could adequately explain the response of legal professionals to AI-enabled legal services (or whether new theory was required).

The second activity was to generate detailed data that would provide the thick description necessary to conceptualise AI-enabled legal services and *explore* its boundaries. As shown earlier, AI remains a contested term, which has frustrated the creation of measures that might allow the use of AI-enabled legal services to be more accurately quantified in the future.

The third activity was *theorising*, an important component of any PhD study. However, having recognised opportunities to theorise would exist, I was aware that theorising about AI was more likely to be abductive in nature (von Krogh, 2018), and that my research design would need to reflect this. Linked to this was a recognition that such an approach to theorising would not generate theory suitable for wider generalisation, given its development within a specific context, which would form part of the theoretical explanation for any observed empirical outcomes. Rather, the objective would be to narrow as far as possible, the range of theories capable of explaining the generated data, and identify those which were most plausible (Bamberger, 2018). A summary of my objectives for the research are detailed below in Table 10. These provided a useful touchstone when developing my research questions and methodological approach.

Table 10. A phenomenon-based research strategy for AI-enabled legal services

Research Strategy	Objectives
Distinguish	<ul style="list-style-type: none"> ● Identify and investigate peculiarities relating to AI-enabled legal services against existing body of knowledge. ● Identify relevant concepts that will enhance the study of AI-enabled legal services. ● Identify inadequacies in theory and knowledge of AI-enabled legal services.
Explore	<ul style="list-style-type: none"> ● Generation of novel data relating to AI-enabled legal services. ● Refine the definition of AI-enabled legal services and generate/refine concepts that can serve as a filter for future data generation.
Theorise	<ul style="list-style-type: none"> ● Generate new theory relating to AI-enabled legal services. ● Expand, adapt, modify existing theory that seeks to explain AI-enabled legal services.

3.3 Theoretical Stance

Having determined an overall strategy for researching the phenomenon of AI-enabled legal services, my next decision was to select a suitable theoretical stance to underpin my research. Management research's theoretical base is drawn from several disciplines, each with its own underlying assumptions, philosophies, and research approaches. While this diversity presents management researchers with several options, it also increases the risk of internally inconsistent research (Saunders, Lewis and Thornhill, 2015).

To help avoid such inconsistency I sought to follow the advice of Crotty (1998) and make my decision-making as transparent as possible. This required me to reflect on the choices I made, and was designed to ensure that my research could be understood and replicated by others. As Van de Ven (2007) argues, clearly showing how I arrived at my findings should make them more persuasive and increase the likelihood that my research impacts both academic and practitioner audiences.

Internal consistency requires alignment throughout a research project. The chosen theoretical perspective and conceptual framework (if one is used a priori) must align with the research problem, purpose of the research and research questions. There is also a need to ensure consistency between field maturity, research questions and the anticipated theoretical contribution (Edmondson and Mcmanus, 2007). Once this has been achieved, a suitable research design and methodology for generating, analysing, and evaluating data can then be selected. Consistency therefore requires researchers to clearly understand and state their ontological, epistemological, axiological, and methodological assumptions (Creswell, 2013).

In seeking to achieve the above, I first set out the key tenets of critical realism, the theoretical standpoint I adopted in my research, explaining how they relate to studying the phenomenon of AI-enabled legal services. I then explain the decision-making process that led to my choice of theoretical position, which compared critical realism with two alternative perspectives (positivism and interpretivism) that are frequently used to investigate technological phenomena.

Theoretical standpoint: Critical realism

Crotty (1998) describes three epistemological positions (what can be known) within the philosophy of science that can underpin a research project. Objectivism posits that an object exists and has independent meaning, beyond the conscious experience of the object by a human subject. In contrast, subjectivism assumes that an object will lack meaning, until that meaning is created by a human. Constructionism offers a third position, arguing that meaning is constructed through the interaction of both the object and subject. Researchers can adopt an epistemological position which best reflects their understanding of the world and use this to guide them in the selection of an appropriate theoretical perspective to underpin their research and choice of research methodology. Internal consistency requires us to highlight the assumptions implicit within a theoretical stance and justify why we are warranted in making them, given the aims of our research (Crotty, 1998).

As with most management domains, the vast majority of research within the field of information systems (IS) has traditionally adopted a positivist theoretical perspective, underpinned by an objectivist epistemology and ontology (Orlikowski and Baroudi, 1991; Williams and Wynn, 2018). Such research assumes the existence of fixed relationships within a phenomenon and utilises hypothesis testing as a means of increasing our predictive understanding of a phenomenon. The theory-led approach constitutes the majority of positivist research, but there is also positivist research which is primarily descriptive in nature. This research lacks a theoretical grounding, with findings simply reported as a 'factual' account of the topic of interest (Orlikowski and Baroudi, 1991).

More recently there has been an increase in the volume of IS research adopting alternative theoretical perspectives, this includes interpretivist studies, with a constructionist epistemology and relativist ontology (Orlikowski and Baroudi, 1991). Interpretivist IS researchers seek to understand IS phenomena through accessing the meanings that individuals ascribe to them. Unlike the aforementioned descriptive research of positivists, interpretivist findings claim only to reflect the research participants' shared understanding of the phenomena, they are not treated as an objective or factual account that can be generalised to other contexts. The aim of interpretive research is, therefore, not to identify covering laws, but to instead understand the deeper structure of the phenomenon, and by

examining social rules and meanings, explain why participants behave the way they do (Orlikowski and Baroudi, 1991).

Critical realism is also gaining traction as a theoretical perspective within IS research (Mingers, Mutch and Willcocks, 2013). Critical realism (Bhaskar, 1997, 1998) is a general research orientation that fuses a realist ontology with a relativist epistemology. It aims to make sense of phenomena by focusing on explaining social processes and events, and through identifying and gaining understanding of the causal mechanisms that underpin them. Critical realism also acknowledges that while individuals possess a degree of agency to shape their social world, they are simultaneously constrained by the wider social structures in which they are embedded. Causality is therefore complex, reflecting the individual's subjective interaction with unseen forces that exert some objective power over them (Houston, 2014).

This interest in causal explanations means critical realist research seeks to move from 'the what' to 'the why' (Vincent and O'Mahoney, 2018), abstracting from empirical events to gain deeper understanding of the pre-structured nature of reality, before returning to an analysis of empirical events, actions and processes in light of this knowledge (Fairclough, 2005). The approach, therefore, helps us to understand the world we inhabit; while simultaneously requiring us to be reflexive about our ethical obligations to others, in light of the knowledge that emerges from research (Houston, 2014; Pilgrim, 2020).

What distinguishes critical realism from positivist and interpretivist positions is its combination of a strongly defined, realist ontology with a more cautious, relativist epistemology, which together allow the use of judgemental rationalism to explain outcomes (Pilgrim, 2020; Wynn and Williams, 2020). Ontological realism means the world and the entities that constitute reality, are considered to exist independently of our ability to perceive or think about them, with reality recognised to have both transitive (changed by the way we talk about and construe the world) and intransitive dimensions (not shaped by our thoughts and actions). Epistemological relativism assumes all knowledge (including scientific knowledge) is socially situated, because our access to reality is mediated by the social structures we belong to. This means all knowledge is value-rich and theoretically informed, hence it is incorrect to conflate the character of the world with what we say

about it. In combination this pair of assumptions allows critical realists to maintain that truth is not relative, but that it remains independent of knowledge. This contrasts with positivist and interpretivist perspectives that commit the epistemic fallacy, whereby the nature of reality (ontology) is conflated with our knowledge of it (epistemology). In positivist research this is reflected in the assumption that what cannot be perceived (and measured) does not exist. In contrast, interpretivist accounts mistake the limitations of our own perceptions and knowledge to be limitations of being itself (Mingers, Mutch and Willcocks, 2013).

Critical realism's assumptions directly impact how researchers develop claims to knowledge, and how the truth of such claims is evaluated against competing knowledge claims. Critical realist research focuses on identifying the causes of past events, not predicting future events. The evaluation of different theoretical explanations is used to determine the most plausible explanation for a phenomenon, in a specific context. While this approach does not allow for the generation of covering laws, the causal explanations identified in one context may be found to exist elsewhere. Hence while broad generalisations are not justified, similar outcomes (referred to as 'demi-regularities') may be observed. Where these are seen, they offer the potential to deepen and generalise (in a limited way) knowledge of causality. Hence, while critical realists argue individual theories can only ever partially capture reality, they also suggest that over time scientific inquiry acts as an error-correction process, selecting those models and theories that most accurately reflect the reality they seek to explain.

The combination of a realist ontology and a relativist epistemology, allows critical realists to employ a stratified representation of the social world. In describing this world and the phenomena found within it, the terms *entity*, *structure*, *event*, *experience* and *mechanism* are frequently employed, hence it is important that their meaning is clearly understood (Easton, 2010).

Entities provide the basic theoretical building blocks of an explanation. Consisting of various forms (physical objects, biological agents, social groups, and rules and practices) they may or may not be directly observed, meaning their existence is inferred through observing their effects. Entities have causal powers and liabilities, which reflect their internal structure, and

which can cause events to occur. In the context of my research, entities included individual legal professionals, law firms and AI software programs.

Entities typically exhibit structure; this refers to the necessary relationships between the components of an entity i.e. its internal objects and practices. The structure of an AI-enabled legal services, can, therefore, include AI software, knowledge databases, individual legal professionals, multi-disciplinary teams and professional norms of behaviour, all of which can affect one another.

Events (or outcomes) are the focus of critical realist research and comprise the external and visible behaviours of people, systems and things. Events are caused by the enactment of one or more mechanisms that are themselves associated with a contingent set of entities and contextual conditions. Critical realism understands social reality as an open system in which both structures and environmental context are subject to change. Events therefore reflect both causal powers (within a social structure) and contextual conditions. This means critical realist research limits its aims to identifying how mechanisms tend to act, in a given context and time. In my research, an important example of an event was the use of AI-enabled services by legal professionals.

Experiences are the sub-set of actual events that it is possible to directly observe or measure, either partially or in their entirety. Hence experiences do not allow us to become directly aware of all the events taking place and the mechanisms that underpin them. For example, observing the use of AI-enabled legal services by legal professionals does not in itself reveal all the events that contributed to the legal professional's decision to use AI-enabled legal services.

Mechanisms are causal forces that make things happen, and which can be identified by their effects. It is through mechanisms we begin to understand how the different entities involved in a causal process are enacted. Mechanisms are understood through first explicating the structure from which they emerge; it is the properties associated with structures and their interaction, that produce mechanisms. The mechanisms themselves are often hidden from us, meaning we infer their existence through experiences we believe them to have caused. When more than one mechanism could potentially explain an event,

critical realism uses judgemental rationalism to identify which theoretical explanation generates the most accurate representation of reality. For example, it is plausible that a legal professional might decide to use AI-enabled legal service based upon a mental calculation of its potential costs and benefits. A competing mechanism might be that the decision to use AI reflects the legal professional following the behavioural norms of their organisation.

In addition to these terms context also plays an important role in developing causal explanations. Context refers to the wider environment within which entities are found. When employing a mechanism-based understandings of causation, we can think of causal factors as what triggers a mechanism to produce an outcome, whereas a contextual condition is merely an enabler of the outcome, it does not do anything active (Beach and Pedersen, 2016).

Together these concepts are used to generate explanations for why an event happened, the fundamental aim of all critical realist research. A simplified representation of this type of causal explanation is provided below (Figure 8), using indicative examples of relevance to this research project. In reality, such formal explanation is not usually feasible due to the complexity of the social world, in which multiple contextual conditions and mechanisms can exist alongside, and interact with, one another.




Figure 8: A causal explanation of AI adoption using critical realist concepts

The above concepts are placed within a stratified ontological framework comprising three nested layers, or domains as they are often referred to. The *empirical* domain consists of what is observed, this reflects an individual’s personal experiences and their perceptions of the events they have witnessed. The *actual* domain refers to where events occur, hence events can still be seen to have occurred even when they are not experienced or witnessed.

The domain of the *real* consists of the structures and causal mechanisms that generate events. This means the reported experience of individuals can only ever be the starting point to understanding what is actually taking place. Critical realists, therefore, seek to generate explanations that are capable of linking personal experiences and events to the underlying mechanisms that caused them.

Using critical realism to research AI-enabled legal services

The critical realist theoretical perspective was supportive of my research strategy to distinguish, explore and theorise about AI-enabled legal services from the perspective of legal professionals. Research into technological phenomena takes place within a complex socio-technical environment (Wynn and Williams, 2020) encompassing: *social structure* made up of different individuals and groups; *material artefacts*, such as AI software programs; and *rules and practices* that govern technology use and shape human-AI relationships. Sitting outside these structures are *discursive entities*, such as language and culture, which can affect the meanings attached to elements of the social structure and causal mechanisms (Archer, 1995; Pawson and Tilley, 1997). Hence, they should be accounted for when researching phenomena such as AI-enabled legal services.

Critical realism ability to analyse such social complexity made it in my view an ideal theoretical perspective for researching AI-enabled legal services, a phenomenon that is still maturing and not yet well understood. For example, a critical realist analysis of user acceptance of AI-enabled legal services, is able to recognise that the outcome of such a process may in part be determined by legal professionals, who possess a degree of agency in their decision-making, while also acknowledging that the decision to use AI is affected by the affordances of the technological artefact, and myriad other mechanisms related to the structure and context that the technology exists within.

The benefits of critical realism became even more apparent when I contrasted it with alternative theoretical standpoints available to me. In seeking covering laws, positivist research overlooks the role of contextual conditions in enabling events (Orlikowski and Baroudi, 1991). This would limit its ability to reveal the full complexity of the social context in which AI-enabled legal services exist, leading to an incomplete understanding of the

phenomenon. The theory-led nature of most positivist research also requires researchers to adopt a predetermined understanding of the phenomenon that aligns to their chosen theoretical model. While this approach makes sense when researching more established phenomena through the use of hypothetical statements, it is less useful when investigating emerging phenomena, such as AI-enabled legal services. Having reviewed the emerging literature on AI-enabled legal services, it was apparent to me that the nature and direction of human-AI relationships is not yet well-understood, meaning I did not want to undertake a theory-driven investigation of the phenomenon. Instead, I wanted to use more exploratory research questions, which would be suited to generating local understanding and knowledge of the phenomenon.

Adopting an interpretivist perspective would have been consistent with the more exploratory research questions I wanted to ask, but it would have required a different set of compromises. Interpretive perspectives by focusing primarily on the perceptions of participants, neglect to examine the wider structural conditions that can explain how and why such experiences and perceptions arise (Orlikowski and Baroudi, 1991). In seeking to explain the behaviour of legal professionals in respect of AI-enabled legal services, I felt it was important that my research strategy could account for the impact of the structure of the UK legal sector and the tensions and contradictions that exist within the professions (Adams, 2020). Additionally, adopting a perspective that denies the separate existence of the technological artifact felt unwise.

In contrast to the above, critical realism provided me with a theoretical perspective that recognised the importance of subjective meanings and context, without denying the separate impact that AI as a technological artefact could have on my research (Smith, 2006). Critical realism, therefore, allowed me to avoid the charge of technological determinism (i.e. that AI is the primary factor shaping social relations and organisational structures); and by adjusting, rather than abandoning, notions of causality, still retain the ability to argue that AI can causally interact with the world, therefore making AI-enabled legal services a worthy research topic.

Researcher reflexivity in adopting a critical realist standpoint

Alongside the aforementioned benefits critical realism offered my research, my decision to adopt a critical realist standpoint was also influenced by my professional and academic background. My experience of working within a law firm environment demonstrated to me first-hand the complexity of such organisations, and the dynamic market context within which they operate, and new technologies are deployed. Adopting a theoretical perspective that embraced the idea of an open systems perspective was, therefore, important to me.

My professional background as an organisational psychologist also attuned me to the fact that the empirical world, accessed through our senses, only provides us with a limited understanding of the processes that explain what is happening in the world around us. For example, processes taking place within the human brain are invisible to the casual observer but can be revealed using technology, such as magnetic resonance imagery, or inferred through observing human behaviour. Hence, I readily related to the stratified ontology of critical realism; and recognised its value in trying to make sense of the acceptance of AI amongst legal professionals.

Working in human resources has also demonstrated to me the socially situated nature of knowledge. I have observed how perceptions of reality vary between individuals, and that there are pluralities of reality experienced by different people exposed to the same phenomenon. For example, when conducting workplace selection processes, it was clear that it was naïve to treat interviews as just a neutral tool for collecting data that would uncover the truth about a candidate's suitability for a job role. In practice different interviewers interpreted candidates' responses in multiple ways, often drawing different conclusions as to the candidate's perceived level of ability. I believe the same issues arise when generating data and trying to interpret the meanings of complex management phenomena, such as AI-enabled legal services. Hence, I share the critical realist stance that data and knowledge are made by people, and multiple realities are an inherent part of the construction of knowledge; but that when combined with a realist ontology it remains possible to go beyond mere description and make cautious suggestions about causes. Having worked as a consultant to professional service firms I am very aware that professionals attach limited value to description, instead they value knowledge and

understanding. I was, therefore, motivated to undertake my research in such a way that my findings would focus on generating causal explanations, with the potential to provide legal professionals with insights that had the potential to change their behaviour in positive ways.

The emancipatory potential of research is important to me and was something I wanted my own research to reflect. Academically I have been influenced by the critical management literature, especially the debate surrounding researcher performativity. Working as a business consultant I have found myself unthinkingly colluding with management decision-makers to further their goals, without reflecting on how my work affected other stakeholders. The impact this can have can be seen through the high incidences of mental health issues in the contemporary legal profession, something I have been a witness too. Most personally, I saw my father who practised law for over forty years alienated from a profession he once loved, as a result of changes in the business model of law firms and the introduction of new technologies. I, therefore, recognise the potential of AI-enabled legal services to have both positive and negative effects on those who interact with them, and am committed to ensuring my research will further the well-being of legal professionals who find AI-enabled legal services part of their professional practice, especially those who may have limited choice in adopting them. Critical realism's ability to accommodate an emancipatory agenda (Danermark, Ekström and Karlsson, 2019; Pilgrim, 2020) therefore appealed to me. This has meant trying to go beyond investigating AI-enabled legal services, simply to learn more about them, to also consider how such services can best be used to generate positive outcomes for different kinds of legal professionals.

3.4 Methodological principles and research design

Having chosen a critical realist standpoint consistent with my overall research goals, I needed to ensure its consistency with my research questions and the methodology used to investigate them. My research questions needed to adhere to critical realist assumptions, by focusing on causal explanations for the outcomes associated with my chosen phenomenon of interest, AI-enabled legal services (Easton, 2010; Wynn and Williams, 2012). When developing my research proposal, the wording of the research questions was kept general and flexible, reflecting the still developing understanding of the phenomenon,

within the professional services field. The questions I generated through my initial literature review were closely aligned to my exploratory research strategy, focusing on *how* AI-enabled legal services have come to be used by legal professionals. The intention being to expose the causal forces within the structural context of UK law firms, and the mechanisms they generate, which together offer a credible explanation for the use of AI-enabled legal services by legal professionals.

My research questions did not make use of existing conceptual frameworks. Instead, I treated my acquired theoretical knowledge as provisional, and did not use a pre-determined theoretical framework to analyse or develop my understanding of the phenomenon (Hoddy, 2019). Instead, my research questions were refined during the process of generating and reviewing my data, during which I identified empirical evidence that raised unexpected questions, which merited further investigation (see chapter 5.2 for further discussion).

In selecting a suitable research design to investigate my research questions, I was guided by a set of five methodological principles that reflect critical realism's ontological and epistemological assumptions (Wynn and Williams, 2012). When applied together these principles enable the development of statements that explain (rather than predict) the events, structures and mechanisms that relate to a phenomenon, and the evaluation of these proposed explanations. The desired outcome is a, "*philosophically acceptable explanation for the existence and operation of a reality that would have logically generated the observable phenomena under examination.*" (Wynn and Williams, 2020, p.52).

These five methodological principles underpinned my research of AI-enabled legal services and subsequently led me to identify a phenomena-driven, critical realist case study as a suitable research design through which to apply them. The five methodological principles are summarised below, followed by an overview of the case study method.

Five methodological principles

Critical realist research should include *thick description of events*, including any actions and outcomes observed by participants, the researcher, or that can be measured empirically. An abstracted sequence of events can then be created and used for further analysis and theory development. Critical realist research, therefore, makes extensive use of qualitative

methods for generating data (although quantitative methods may be used in addition to this). In my research, in order to understand the use of AI-enabled legal services, I selected a methodology capable of generating detailed descriptions of: an AI-enabled legal service; its associated outcomes; and the wider context it sat within.

I also sought to overcome the perceptual limitations of humans, by *generating data from multiple sources, using different methods*. Rather than relying solely on interviewing legal professionals to develop my understanding of AI-enabled legal services, each participant was also asked to create a narrative timeline of their experiences of AI-enabled legal services in advance of the interview. In addition to this I also conducted background discussions with organisational sponsors and received demonstrations of the AI-enabled legal services in use at each firm, which allowed me to observe them first hand (these methods are discussed in more detail in Chapter 4).

Critical realist research needs to *identify the structures and contextual conditions* that are causally relevant to an outcome, and the relationships amongst them. This helps transform individual participant accounts into a more abstract theoretical perspective. As it is not feasible to consider every aspect of a complex socio-technical environment, descriptions need to focus on those aspects most relevant to the research questions. When analysing the data generated in this research, I used a thematic analysis and employed process tracing (see Chapter 5 for further details) to identify the structures and conditions that impacted both how legal professionals came to understand and accept the use of AI-enabled legal services.

Both types of analysis employed, *Retroduction*, a key tool of critical realist research that involves asking of the observed phenomenon, *'what must be true for this to be the case?'* (Oliver, 2012 p.379). Retroduction leads to the identification of one or more plausible causal mechanisms, capable of explaining the outcome of interest, which can then be compared to extant theories to identify areas of commonality and ensure a parsimonious theoretical explanation.

These explanations could then be *tested against the available empirical data* (a combination of my observations and reported participant experiences) to ensure the proposed causal

mechanisms adequately reflected reality. With a mechanism’s explanatory power determined by the extent to which it was corroborated by the available data.

In this research, where competing explanations for an aspect of AI-enabled legal services emerged, a mechanism that demonstrated explanatory value for multiple individuals working in different law firms was deemed more powerful than an explanation whose power was limited to a narrower sub-set of legal professionals within a single context.

Research design: Critical realist case study

The case study approach is ultimately grounded in the in-depth investigation of one or a small number of instances of a contemporary phenomenon, within its real-life context. The aim being to make sense of the phenomenon, rather than the field within which it is situated. The approach is used across disciplines and is in widespread use in the Management domain amongst qualitative researchers (Plakoyiannaki, 2020). Despite this, the term case study defies definition, in large part because case studies appear across different research paradigms, with varying guidance for their use (Fletcher *et al.*, 2018). The distinctiveness of critical realist case studies becomes more apparent through a comparison with both positivist and interpretivist case studies (summarised in Table 11).

Table 11: Different theoretical standpoints in case-based research

Positivism	<ul style="list-style-type: none"> • Causes are narrowly conceptualised • Research focuses on effects of particular causes across cases • Parsimonious causal theories are favoured • Generalisation across cases attempted • Causal conditions in focus, and if mechanisms theorised, they are viewed in a similar way to intervening variables
Interpretivism	<ul style="list-style-type: none"> • Focus on description of a phenomenon in its local context • Underlying epistemological and ontological assumptions not compatible with causal case study methods • Generalisation across cases not attempted
Critical Realism	<ul style="list-style-type: none"> • Causes are broader and more complex, often working in conjunction • Research focuses on complex causal patterns • Generalisation only to small, bounded populations • Singular causation also in focus • Mechanisms are at the core of research

Source: Adapted from Beach and Pedersen, (2016)

In positivist case study research (e.g. Eisenhardt and Graebner, 2007; Yin, 2018) the aim is to collect detailed qualitative data as a means to help generate theory about a phenomenon that is not yet well-researched and understood. The underlying assumptions of this approach reflect a scientific orientation to research, whereby empirical data is used to generate theoretical frameworks and hypotheses that can then be tested in other contexts to assess their potential for generalization across settings (Leppäaho, Plakoyiannaki and Dimitratos, 2016). In this sense the positivist case study can be seen as a preliminary step prior to the 'real research' taking place (Plakoyiannaki, 2022).

Interpretivist case study researchers (e.g. Walsham, 1995) regard reality as being socially constructed, with knowledge of the social world dependent upon human interpretation. Case studies offer interpretivists a means to generate data and develop their understanding of a phenomenon, through recognising its complexity and the way in which it interacts with its context. The findings of an interpretivist case study, therefore, reflects the researcher's interpretation of the phenomenon (Plakoyiannaki, 2022). Insights from such research are therefore localised to the specific context, with case studies not regarded as a method to produce generalisable propositions that can be replicated elsewhere (Leppäaho, Plakoyiannaki and Dimitratos, 2016).

Critical realist case studies (CRCS) (e.g. Easton, 2010), are used to explore the interactions between events, structure and context with a view to identifying underlying mechanisms that can explain empirical happenings linked to the focal phenomenon (Avenier and Thomas, 2015). This means CRCS tend to focus on sets of causes and causal complexity, in addition to singular causal claims. CRCS do not seek to identify causal explanations that function at a general level, but instead focus on explanations that are context based. This means generalisations of causality are more limited in scope and typically only made to small, bounded populations of similar cases (Beach and Pedersen, 2016). This focus on 'deep' causal mechanisms that cannot always be measured directly, means critical realist researchers are required to generate explanations that go beyond simple description; the aim being to overcome the tension between context and causal explanation that exists between positivist and interpretivist case study research (Leppäaho, Plakoyiannaki and Dimitratos, 2016). Avenier and Thomas (2015) describe this as 'abductive explanatory'.

While not the only option available to critical realist researchers, case studies offer scope to realise the five methodological principles outlined previously. This has meant “*several critical realist researchers have identified the case study method as the best approach to explore the interaction of structure, events, actions and context to identify and explicate causal mechanisms*” (Wynn and Williams, 2012, p795). Reflecting this viewpoint is the increased use of CRCS to investigate complex organisational phenomena, such as IT-led organisational change and how information systems evolve over time (Allen *et al.*, 2013; Henfridsson and Bygstad, 2013; Volkoff and Strong, 2013). This growing consensus lent further support to my decision to use a case study design for investigating the phenomenon of AI-enabled legal services, within the specific context of the UK legal sector. The precise design of the case study and the methods chosen to generate data are discussed in the next chapter.

4. CASE STUDY DESIGN AND METHODOLOGY

This chapter details the case study design used in the research and the specific methods chosen to generate data. First, I set out the assumptions that informed the design of my case study, explaining their relevance to the overall process of data generation and analysis. I then explain the process through which cases were selected, organisational partners were identified, and participants selected. Descriptions of both organisational partners are provided for background context. I then discuss the impact of wider contextual factors that impacted the methods choice – undertaking research during the COVID pandemic; conducting research with members of a professional elite; and research ethics. In light of the above, I then outline the chosen methods for generating data: observation; reflective exercise; and electronic semi-structured interview.

4.1 Assumptions in the case study design

My case study is underpinned by several assumptions relating to how causality is understood, together these ensure consistency with the critical realist perspective. These assumptions differ from those that underpin case studies from different philosophical perspectives. These first two assumptions – ontological determinism and causal asymmetry – reflect the epistemology and ontology of critical realism. The final assumption reflects critical realism's focus on understanding causation through the identification of mechanisms, as opposed to the use of counterfactuals.

Ontological determinism means when analysing a case, all hypothesised causal relationships are seen in categorical terms; this means the relationship between a cause and an outcome either exists or it does not, we are not dealing with within-case probabilities. For example, in my research it is not possible for a legal professional to simultaneously be both a user and non-user of AI-enabled legal services. Epistemologically, however, case studies do still deal in probabilities, as our knowledge of the world can never be certain. This means when seeking to explain causal relationships, our confidence will reflect the quality of the available evidence. Hence, when the available empirical evidence does not support our

existing theoretical explanations, we need to revisit our theory and seek an explanation for why the anticipated relationship did not occur (Beach and Pedersen, 2016).

Asymmetric causality means explanations about the causes of an outcome do not explain what causes the absence of an outcome. Instead, we assume that the causes of an outcome and its negative will be very different (Beach and Pedersen, 2016). For example, in my research the causal mechanisms identified during within-case analysis of a single legal professional, and which explain why that individual used AI-enabled legal services, does not indicate which mechanisms explain why legal professionals do not use AI-enabled legal services. The assumption also means that when conducting cross-case analysis any generalisations are restricted to a population of cases with similar causal conditions and the same (positive) outcome. For example, in my research I did not seek to generalise my findings to legal professionals in very different contexts, for example those working in other legal jurisdictions, or law firms with radically different organisational models.

Causality can be understood in different ways by researchers. In my research I treat causality as a mechanism, unlike most positivist research in which counterfactual accounts of causation are the norm. While the counterfactual approach provides a logical foundation for making claims about causation based on correlations between variables, it does not offer insight as to *how* the identified cause contributes to producing the outcome (Beach and Pedersen, 2016). In contrast, mechanism accounts of causation focus on tracing causal processes in an actual case, to reveal evidence of the 'connection' (mechanism) linking a cause and an outcome at the within-case level. Mechanism accounts can also account for instances where more than one cause may be implicated in an outcome, by evaluating the theoretical uniqueness of the available empirical evidence, to determine which cause(s) is(are) most plausibly linked to the outcome.

It is also important to recognise that mechanisms can be understood in two broad ways, which Beach and Pedersen (2016) term *minimalist* and *systems* understandings. In minimalist understandings the causal process between a cause and an outcome is acknowledged but not unpacked theoretically in any detail, which weakens any causal claims that are made. The systems understanding, which seeks to provide a detailed explanation of causal mechanisms, more accurately reflects the critical realist position,

where the aim is to develop a theoretical understanding of the mechanism that links a cause and an outcome. The *system* label is used as mechanisms are understood to be complex and likely to comprise several different elements. While the minimalist understanding might at first glance appear inferior to the system understanding both can be useful (Beach and Pedersen, 2016). In my research I utilised both understandings, reflecting the type of analysis taking place and the availability of empirical data, being careful to ensure that the claims I made were justified given the chosen understanding.

The methodological implications arising from these assumptions are summarised below (Table 12); their impact on my research is referenced in my discussion of case selection (see section 4.2) and in relation to the process of data analysis (Chapter 5).

Table 12: Methodological implications of causal case study research

<p>Case selection</p> <ul style="list-style-type: none"> • Cases selection requires causal relationships to be categorical in nature. • The case population needs to demonstrate causal homogeneity in order for (limited) generalisation to be possible. Causal analysis should focus on typical cases, within the overall case population, provided there is sufficient empirical material available.
<p>Generalising potential</p> <ul style="list-style-type: none"> • To maximise the potential to generalise mechanisms beyond the case they were identified in, a bounded population should be selected from the overall set of cases.
<p>Concept definition</p> <ul style="list-style-type: none"> • When tracing causes in a case, concepts should be defined in categorical terms to help minimise heterogeneity between cases.
<p>Generation of causal explanations</p> <ul style="list-style-type: none"> • System explanations require process tracing as the method of analysis, as this mode of analysis is capable of identifying explicit causal mechanism that can link a causal factor to an outcome. Where process tracing is not feasible, only minimalist causal explanations can be generated. • Causal claims are asymmetric and do not allow claims to be made about what happens when a causal factor is not present.

4.2 Case study design

Case selection – unit of analysis

When employing the case method, researchers need to be transparent about case selection (Plakoyiannaki, 2020). Case selection should reflect the unit of analysis – the focal entity being investigated (i.e. the ‘what’ or the ‘whom’ of the research); this is important as it determines what the researcher can make claims in relation to. Case selection also has implications for data generation and analysis and how findings can be reported (Fletcher and Plakoyiannaki, 2011). In critical realist research we are focused on understanding the mechanisms that link a cause (or set of causes) with an outcome. The case is, therefore, the unit in which the causal process plays out (Beach and Pedersen, 2016), from the occurrence of the cause to the anticipated outcome.

In my own research while AI-enabled legal services was the phenomenon of interest, the causal processes linked to it were played out at the individual level, for example, did a legal professional use AI or not. Hence the findings of my research were also reported at this level, with each participant in the research treated as a separate case. The empirical unit of observation that was used to generate data was also individuals, with the case population comprising UK legal professionals currently using AI-enabled legal services.

Case selection strategy

Once the unit of analysis has been determined, a strategy for selecting cases can be chosen. Fletcher *et al.* (2018) highlight two contrasting strategies – theory-driven and phenomenon-driven case selection. Theory-driven cases are chosen prior to data generation, because of their potential to extend or revise extant theoretical relationships between constructs. This approach is more appropriate when understanding of the focal phenomenon has reached maturity. In phenomena-driven case research there is not an a priori definition of the case as understanding of the phenomenon and its context comes through the research process rather than the extant literature. Cases are, therefore, identified in light of the data that is generated, with selection reflecting the extent to which the case allows the phenomenon of interest to be captured, documented, and conceptualised. This approach is common

amongst interpretivist case study researchers and is better suited to researching phenomena in the embryonic or growth phases of research.

Critical realist case studies focus on developing detailed causal explanations for a small number of events in a specific setting. This means case selection is not primarily based on generalising potential, instead there is a focus on cases where empirical events are representative of the phenomenon of interest. For example, Henfridsson and Bygstad (2013) chose to study the evolution of information systems in a company regarded as being prototypical of successful IS evolution. This maximised the likelihood of them being able to identify concepts and generate theory capable of explaining how IS successfully evolve (but would offer limited insight into why IS evolution might fail).

The immaturity of AI-enabled legal services as a phenomenon led me to adopt a phenomenon-driven strategy of case selection. This meant focusing on empirical settings in which AI-enabled legal services were being used by legal professionals. I, therefore, chose law firms purposively, to maximise the likelihood that they would contain a sufficient number of detailed cases to develop an understanding of the relationships that develop between legal professionals and AI-enabled legal services. The case population was not determined prior to data generation; instead, the process was dynamic, with my decision on where to focus data generation guided by the process of analysis, as the causal relationships of interest became apparent (see Chapter 5 for a detailed discussion).

Sampling strategy

Sampling is complex in case study research as the choice of sample needs to align to the overall research strategy and case study design (Fletcher and Plakoyiannaki, 2011). Phenomenon-driven cases employing qualitative methods focus on purposeful sampling strategies, which are designed to identify appropriate and information-rich cases. This contrasts with representative and random/probability sampling, which is often seen in quantitative research (Fletcher *et al.*, 2018).

There are several different approaches to purposeful sampling (see Patton, 2014 for a summary). In my research I initially employed a selective sampling strategy (Sandelowski,

1995), whereby I created a justifiable set of criteria to help me identify where it might be possible to research the phenomenon of AI-enabled legal services (Table 13).

Table 13: Sampling criteria for the study of AI-enabled legal services

Criterion	Justification
Established use of AI-enabled legal services	To maximise the potential size of the case population and increase the likelihood that any identified cases would be information rich.
UK legal sector context	To increase population homogeneity, the use of AI-enabled legal services was limited to the UK offices of commercial law firms. To simplify the process of data generation from an ethical and data processing perspective, law firms needed to be headquartered in a UK legal jurisdiction.
Access to legal professionals with experience of using AI-enabled legal services	Access is a significant challenge when conducting research in the field of professional service firms. No pre-existing datasets were identified, meaning data generation was required for analysis to take place. Access at the individual level was necessary, given the anticipated unit of analysis.

Based on the above criteria a list of potential organisational partners was created. Eleven UK law firms that satisfied the first two criteria were identified and contacted. Of these, two firms were willing to provide the necessary organisational access. An anonymised description of each firm and the AI-enabled legal services it employed is detailed below. To preserve the anonymity of the firms, some distortion of the facts and figures used to describe each firm has been necessary; I also describe the types of Artificial Intelligence in general terms without reference to specific processes or software packages.

National LLP

National LLP, from now on referred to as National, is a leading commercial law firm, with a turnover in excess of £100m and over 1000 people working across its network of UK offices. This places it amongst the top 50 UK headquartered firms by turnover. Like the majority of large UK law firms, National has a Limited Liability Partnership (LLP) ownership model. This means the firm is owned and run by a group of partners who comprise a minority of the people who work for the firm. Unlike a traditional partnership, in an LLP partner liability for the firm is limited in a similar way to owners of a limited company.

National provides a wide range of legal services to organisations in both the private and public sectors, and to private individuals. The firm undertakes two distinct types of legal work – claims-related and transactional. Lawyers working on claims-related work focus on representing defendants, for example a company that has had a claim for personal injury brought against it by one of its employees. The firm handles a high volume of claims-related cases, many of which follow a well-understood process. Typically, the firm charges a fixed fee for this type of legal work, irrespective of how long the work takes. The firm's transactional departments cover several different legal disciplines including corporate, employment and real estate. When providing legal advice to clients, transactional teams more typically charge for their services using the billable hour model, in which a client pays for the amount of time spent undertaking their work.

National makes extensive use of Legaltech, including AI, when delivering legal services to clients. The firm's use of technology has seen it win industry awards. The types of AI in use at the firm includes expert systems, supervised machine learning and natural language processing. These technologies have been developed through collaboration with commercial and non-commercial partners. While certain tools have been developed from scratch, meaning they are only available to National, the majority are based on software that is commercially available, but which required adaptation in order to integrate it into National's own processes. This has resulted in AI-enabled legal services being implemented across several different practice areas, with their use most prevalent in claims-related work. These processes support the work of fee-earners and can generate data insights, provide decision-making support, ensure compliance, and automate basic administrative processes.

The development of AI-enabled legal services has been coordinated through a dedicated innovation team, responsible for procuring, developing and deploying the technology that underpins the processes as well as providing training to fee earners. The innovation team contains permanent members from a variety of professional backgrounds, this includes qualified solicitors, IT specialists and project managers. The team also has fee-earning solicitors seconded to it on a project-by project basis, to support the development of AI-enabled services, which relate to the individual fee-earner's area of expertise. Secondees typically maintain their fee-earning duties, devoting only a minority of their time to the work of the innovation team. There is also the opportunity for trainee solicitors and solicitor apprentices to undertake a six-month placement in the team as part of their professional training, during which they work exclusively for the innovation team.

Global LLP

Global LLP, from now on referred to as Global, is a leading commercial law firm, with a turnover in excess of £500m. The firm has offices across the world employing over 1,000 qualified lawyers and 3,500 people in total; the majority of whom are based in the UK. This places it amongst the top 20 UK headquartered firms by turnover. Global seeks to offer its clients a comprehensive service, meaning it has lawyers working across a wide range of legal specialisms, of both a contentious and non-contentious nature. Contentious work is focused on resolving client disputes, in contrast non-contentious work involves providing advice to clients and managing the legal aspects of a transaction e.g. the buying or selling of a company. Global's size and reach means the size and value of the work it undertakes for clients (sometimes referred to as 'matters') is on average significantly higher than National. More recently, Global has diversified and now offers clients complementary professional services in addition to its core legal services. Several of the firm's departments receive global rankings for the quality of their work.

Global has digitalised the vast majority of the firm's processes; many processes now feature elements of automation and the use AI. Global has been using AI significantly longer than most UK law firms (including National), and while the types of AI in use at the firm is similar to many other firms (expert systems, supervised machine learning and natural language processing) the sophistication of these tools is often much greater. Global regards itself as

‘problem’ rather than ‘technology-led’ meaning the development of its legal processes (and its use of AI) is driven by the needs of clients, rather than the capability of the technology. This has meant while Global does collaborate with external commercial partners, it has significant internal capacity to develop its own software; many of the firm’s most effective and widespread tools have been developed internally. In certain specialisms, Global’s use of AI, combined with its large scale (relative to most law firms), has meant it is one of a small number of law firms capable of undertaking high-volume, transactional legal work that cannot be delivered cost-effectively using traditional lawyer-centred processes. This has seen Global develop into one of the leading users of technology within the legal sector, something which has been widely recognised through industry awards.

Global has a centralised ‘solutions’ team, responsible for procuring and developing technology for the firm. The firm also had a dedicated ‘delivery’ team that is responsible for working directly with each of the firm’s different departments, to identify innovative ways to meet the needs of clients. Where technology is a core element of the innovation, the team works in collaboration with the solutions team. The delivery team is currently expanding, so that each legal specialism has a dedicated member of the delivery team working with them. Delivery team members typically come from a fee-earning background, but once they join the delivery team they work for it on a full-time basis and are no longer required to undertake fee-earning duties. The combined solutions team and delivery team is significantly larger than the innovation team at National.

Identification and recruitment of participants

My strategy for identifying and recruiting participants was developed in partnership with the organisational sponsors of my research at National and Global. Negotiating access to participants within professional services firms is acknowledged to be difficult, with specific challenges arising from the confidential nature of the work being conducted, the intensive nature of the work-environment, and the diffused nature of decision-making, which means access often has to be renegotiated at multiple levels (Karjalainen, Niemistö and Hearn, 2015). Once organisational access to potential participants has been secured, ‘cognitive access’ to individual professionals is dependent upon gaining their acceptance and consent

(Symon and Cassell, 2012). This requires rapport building and the development of trust, which can be a time-consuming process (Empson, 2018).

My strategy for identifying and recruiting participants was, therefore, pragmatic and opportunistic, as while I was aware that the composition of the participant population would affect my opportunities for data analysis and the claims I would be able to make in my research; I needed to balance my aspirations against the reality that many of the legal professionals, who were my potential participants, would not be interested in participating, despite the support of my organisational sponsors (Karjalainen, Niemistö and Hearn, 2015; Petintseva, Faria and Eski, 2019). To maximise the likelihood that a suitable participant sample was secured, I drafted a collaboration agreement that described the respective responsibilities of myself and the organisation sponsor. By detailing the tasks each party would undertake and the timescales involved, I sought to manage my sponsors' expectations about the work involved, and provide them with an opportunity to challenge the assumptions that underpinned my proposed sampling strategy and methods of data generation (Symon and Cassell, 2012). The signing of the collaboration agreement by my sponsors at both National and Global (neither of whom asked for any amends to be made) gave me confidence that my approach to participant recruitment was realistic.

Once I had reached an agreement with my sponsors, I followed the advice of Karjalainen, Niemistö and Hearn (2015) and sought to identify and secure participants through multiple access points, to increase the likelihood that suitable participants would be made aware of the research and volunteer to take part, given no accurate record of the population of legal professionals using AI-enabled legal services existed in either law firm. Utilising the knowledge and professional networks of my organisational sponsors a list of potential participants was generated, all of whom were known to have experience using AI-enabled legal services. I chose not to restrict the sample to participants *currently* using AI-enabled legal services, as I was interested in understanding the experiences of both adopters and non-adopters of AI-enabled legal services. Beyond this I judged that no further selection criteria were necessary. My rationale was that probability sampling was not required to answer my research questions; instead my priority was maximising the size of the sample frame of individuals capable of offering insights about their experience of AI-enabled legal

services. My selection of organisational partners had also already ensured a high degree of homogeneity amongst potential participants, given all were legal professionals, working within large commercial UK law firms, which would support my process of analysis. I therefore adopted a self-selection sampling technique, with all individuals who were known by the organisational sponsor to meet the inclusion criterion, approached by email with an invitation to take part in the research (See Appendix 2). The email was sent by the organisational sponsor using their organisational email address, as this would be more likely to ensure the email was read than if it was received from an unknown source (i.e. myself as the researcher). Those recipients interested in taking part in the research were directed to respond to me by email (rather than the sponsor) in order to preserve their anonymity.

As the focus of my research was to better understand the phenomenon of AI-enabled legal services, sampling was designed to generate information rich data, rather than representative opinions (Petintseva, Faria and Eski, 2019). To maximise the likelihood of information rich cases that would be suitable for analysis using methods such as process tracing (Beach and Pedersen, 2016), I asked my sponsors to highlight individuals whose role and experience meant they were highly knowledgeable of the firm's use of AI-enabled legal services. In addition to receiving the aforementioned email inviting them to participate, these individuals received a further personalised invitation from the project sponsor encouraging them to participate in the research, in the event that they did not respond to the initial email.

Recognising the difficulty in accessing participants solely through my sponsors (who had limited time available to support my research and whose knowledge of the overall population was imperfect), I also secured permission to directly approach members of National and Global who were already part of my professional network, as a result of my previous work within the legal sector. I approached all such individuals by email inviting them to take part in the research. As with all other potential participants, they were required to meet the inclusion criterion of having experience using AI-enabled legal services.

A final technique I used to secure volunteers was snowball sampling, which is recognised as a useful technique for identifying participants amongst an overall population that is hard to identify and/or difficult to access (Symon and Cassell, 2012; Petintseva, Faria and Eski,

2019). After I completed each interview, I asked my participants if they were aware of any colleagues who could make a valuable contribution to the research, due to their experience of using AI-enabled legal services. While snowball sampling typically leads participants to volunteer potential participants who are similar to themselves (which can promote a high-degree of homogeneity in the data generated), I sought to mitigate this by highlighting to participants that I was interested in participants whose views might be similar or different to their own, so long as they had personal experience of using AI-enabled legal services. I also stressed that their anonymity would be preserved and that I would not mention their involvement in the research when approaching anyone they highlighted to me.

Sample size

The number of participants required to undertake high-quality interview-based research is dependent upon a combination of epistemological, methodological and practical factors (Baker and Edwards, 2012); with any numerical indication given at the proposal stage of a project, merely for guidance purposes, rather than something that can be precisely calculated (Saunders and Townsend, 2016). Once data generation is underway, data saturation (the point at which nothing new is apparent in the additional data generated) is frequently used in qualitative research to determine when data generation can be ended (Petintseva, Faria and Eski, 2019).

The thematic analysis I planned to undertake with my generated data was compatible with the larger sample sizes (25-30 participants) recommended for general qualitative studies (Creswell, 2013); although data saturation typically is seen to occur after a smaller number of interviews (9-17) have been analysed, suggesting smaller samples can be sufficient (Hennink and Kaiser, 2022). In contrast, my plan to identify causal pathways using process tracing (Beach and Pedersen, 2016), would involve analysing a single case in detail, before then assessing whether any causal processes could be generalised to a small, bounded population. Guest, Bunce and Johnson (2006) suggest that when dealing with homogenous populations, a sample size of 12 is usually sufficient.

The exploratory nature of my research meant I could not predict in advance whether my analysis would benefit from further comparative analyses amongst specific participant

groups, within the overall sample population; for example, comparing participants working for National to those working for Global. Securing a sample in each firm that was large enough to facilitate such comparisons (following Guest, Bunce and Johnson, 12 participants per group) was, therefore, an additional consideration when estimating my required sample size.

Based on all the above, and leaving myself a degree of flexibility to overcome issues such as participants withdrawing from the research, I agreed with each organisational sponsor that we would aim to secure a minimum of 15 volunteers in each firm. Both sponsors agreed to this, indicating that this number of participants was feasible given their estimated size of the population who met the inclusion criterion in their firm.

Following the above steps led to 21 volunteers who met the inclusion criterion agreeing to participate in the research, 8 of whom worked for National and 13 for Global. This meant the final sample size was below the figure I had originally estimated, despite each sponsor inviting potential participants to volunteer on more than one occasion. On reflection, the lower than anticipated number of volunteers, potentially reflected both the high-value experts attach to their time as a measure of productivity (Stephens, 2007), and the additional demands placed on professionals during the Pandemic (Rahman *et al.*, 2021). A failure to account for the interaction of these factors led to both me and the organisational sponsors being overly optimistic about the expected participant response rate.

Chronologically data generation at National took place between November 2021 and July 2022; data generation at Global was undertaken between September and November 2022. In terms of the overall sample, I judged data saturation to have been achieved, as the final interviews I conducted (which took place in late October at Global) led to limited new data being generated, and no new names being suggested for the purposes of snowball sampling. I, therefore, felt that the overall sample was large enough for me to undertake my planned analysis using the population as a whole. Amongst the 8 participants at National, the final interviews I conducted there in late June 2022 continued to generate additional insights. Although further participants were sought beyond this date, I was not able to secure additional volunteers, meaning saturation was not achieved. This meant my ability to make comparisons between participants at National and Global was compromised.

Selection of case population

For the thematic analysis data from all 21 participants was used. This reflected my objective of using the thematic analysis to answer the more general research question, '*How are AI-enabled legal services understood by the legal professionals using them?*' In answering this question, the data generated from all participants was of relevance, as I wanted to capture in as much detail as possible the range of views held amongst a relatively homogenous population of legal professionals.

Theory building process tracing takes place within a single case. The selected case is chosen from amongst a bounded set of cases, each of which contains a similar set of plausible causal factors linked to the outcome of interest. In my research two bounded populations were identified; the members of each bounded set shared a set of potential causal factors, which when taken as a whole, made them distinct from the members of the other set. Within each bounded set a typical case was then identified for process tracing analysis. The process for the selection of typical cases is discussed in Chapter 5).

4.3 Wider methodological considerations

It is important to recognise that the selection of specific methods for data generation needs to reflect both theoretical concerns and contingent events. In this section I outline wider factors that shaped my final selection of methods, these include the impact of the COVID-19 Pandemic on academic research; conducting research within elite professional services firms; and research ethics.

My decision to use the term *data generation*, rather than the more common *data collection*, is designed to acknowledge that the data in my research would not exist if it were not for my involvement as a researcher, for example, asking questions in an interview, or providing participants with a reflective exercise to complete. This is in contrast to naturally occurring data that is collected by researchers, but which was generated without their direct involvement e.g. a policy document written for employees by their organisation (Paulus, Lester and Dempster, 2013).

Impact of COVID-19 Pandemic on the research

Having started my PhD research in October 2019, I was less than six months into my studies when the first UK government coronavirus lockdown began in March 2020. Social distancing measures, which remained until July 2021, placed heavy restrictions on people's everyday lives and the type of academic research that was permitted during this period (Institute for Government, 2022). This meant research projects that were already in progress required significant adaptation to continue, and many projects were forced to halt completely. This section outlines the changes I made to my own research, the rationale for my decisions at the time, and my reflections on this experience, having now completed the process of data generation.

In 2020, when making these decisions, it was highly uncertain what path the pandemic would take, and as a corollary to this, the likely severity and duration of the government measures that would be needed to control it. This meant in addition to ensuring that my chosen methods were consistent with my broader research design, I recognised that my chosen methods needed to be resilient and flexible enough to respond to any disruptions that might arise in the future (Rahman *et al.*, 2021). It was important to consider the resilience of methods from both my perspective as the researcher and the participants involved in the research. My guiding principles were, therefore, that my methods needed to be flexible enough to adapt to both changes in my research environment and my participants' occupational context.

In comparison to many researchers, I was in the relatively fortunate position of being able to adjust my research design, to make it compliant with government and university guidelines, prior to starting the process of data generation. Despite this, several challenges remained. Government measures were subject to frequent revision, which made long-term planning difficult. Potential law firm partners were experiencing significant changes to their business operations. This included an unplanned shift to home working for the vast majority of legal professionals, and where this was not feasible, making workplaces COVID-compliant. Such rapid change and the accompanying pressure put on organisational resources meant organisational gatekeepers were reticent to enter into new research projects. Peer-reviewed guidance on COVID-compliant methods for data generation was

also less readily available than at the time of writing. While such papers did begin to appear in the second half of 2020 (see Lobe et al, 2020; Silverman, 2020), much relevant research (for example, Rahman *et al.*, 2021; Tremblay *et al.*, 2021) was not published until 2021. Most of my key decisions relating to research design were made in early 2021, which meant that while they were informed by published findings on the use of remote qualitative methods, it was not clear how generalisable such approaches would be to the specific context of research during a pandemic.

Although my methods needed to change, my overall research strategy remained focused on using myself as an instrument of data generation, to understand legal professionals' experiences of AI-enabled legal services. To determine the most appropriate methods for data generation I first identified the immediate challenges of undertaking research during the COVID pandemic. Physical distancing and time constraints were two factors that had the potential to either prevent the research from taking place, or if it were to go ahead, undermine its quality (Tremblay *et al.*, 2021). Pre-pandemic I was considering undertaking ethnographic research within a law firm environment. This would have offered me the opportunity to understand and analyse how legal professionals interact with AI-enabled legal services as a function of their daily activities, rather than trying to gain insight through asking participants to respond to my research questions. Spending extended periods of time within an organisation could also have helped me to develop the high levels of rapport and trust with participants, which are regarded as an essential feature of high-quality qualitative research (Denzin and Lincoln, 2017), and of particular relevance when trying to access elite professions (Empson, 2018). However, so long as social distancing remained in place, such an immersive research approach was not feasible.

While pausing my research until all physical distancing restrictions were relaxed would have overcome this issue, the impossibility of predicting when this might happen, when combined with the time constraints of my PhD studies, meant this was an unrealistic option. I therefore chose to pursue the use of alternative qualitative methods that would allow me to explore AI-enabled legal services as a workplace phenomenon, while also maintaining physical distance from participants. This led me to identify electronic interviewing (Salmons, 2014) as the primary method for generating data via synchronous computer-

mediated interactions between myself and participants. I also identified participant-generated narrative timelines (Kempster and Parry, 2011) as a complementary and asynchronous method of generating data that would offer participants flexibility both in terms of when and where data could be generated (Tremblay *et al.*, 2021).

The adoption of electronic interviews, while overcoming the challenge of social distancing, raised potential challenges of their own. First, I considered the impact that electronic interviews can have on the quality of the data generated. Historically, in-person interviewing was commonly regarded as the interview 'gold standard', due to the belief that in-person interviewing (when compared to alternative interview methods) generated superior levels of rapport between the participant and researcher, reduced the risk of participant over-disclosure, and allowed researchers to more easily 'read' body language and non-verbal cues (see Jenner and Myers, 2019 for a detailed review of the literature). More recently conducted research has, however, begun to challenge these assumptions, reflecting the fact that online encounters are now just as 'natural' for many people as face-to-face conversations (Lobe, Morgan and Hoffman, 2020; Silverman, 2020). Indeed, while the process of creating and maintaining rapport varies depending on the interview format, there is evidence that rapport can be established quickly in electronic interviews when email correspondence takes place prior to the interview (Gray *et al.*, 2020). The functionality of interview software also impacts rapport building, with more recently developed platforms such as Zoom (Zoom Video Communications Inc, no date), reported to play a positive role in forming and maintaining rapport by over two thirds of participants (Archibald *et al.*, 2019). Hence, there is no reason to assume that data generated in-person is any more 'authentic' than electronically generated data (Silverman, 2020); indeed, data quality was found to be comparable (in research using both face-to-face and electronic interviews), even when discussing sensitive topics (Jenner and Myers, 2019). This gave me confidence that using electronic interviews to research the impact of AI-enabled legal services (another technology-mediated process) should not have a detrimental impact on the quality of the data generated.

On a more practical level, I also had to consider whether legal professionals would be willing and able to interview electronically, rather than face-to-face. Scheduling interviews with

busy professionals during a Pandemic was a likely challenge (Rahman *et al.*, 2021; Tremblay *et al.*, 2021), meaning the low-cost and flexibility of electronic interviews would be advantageous (Archibald *et al.*, 2019). With the majority of legal professionals adopting hybrid working patterns during the Pandemic, the ability to undertake an electronic interview from a location determined by the participant, reflected the broader ways in which legal professionals were conducting their work with clients and colleagues. I also judged that it would help legal professionals minimise the likelihood of distraction during the interview, and help them maintain their anonymity (Archibald *et al.*, 2019).

In terms of selecting a specific software package, in 2020 the accessibility and functionality offered by Zoom for electronic interviews was evaluated to be superior to other available platforms (Lobe, Morgan and Hoffman, 2020); it was also compliant with the University of Sheffield data security policy, which I was required to adhere to. Having already used Zoom extensively in my professional work with law firms, I was aware that the technology would be acceptable to organisational partners (from an organisational risk-perspective), compatible with the computer hardware used by law firms, and that its functionality would be familiar to legal professionals. Together, these reasons suggested Zoom would be an ideal platform through which to conduct my research.

Data generation within an elite profession

Having worked extensively with legal professionals, I was aware that generating data within elite professional services firms would pose distinct methodological challenges (Empson, 2018). Elites typically derive their status through their privileged access to knowledge, institutions, and money (Odendahl and Shaw, 2002). Within the context of my research, all of my potential participants could be considered members of an *elite*; reflecting a combination of the high standing that Law enjoys as a profession in the United Kingdom (UK), and the leading status of Global and National as law firms within the UK legal sector. The reputation and role of individual participants within Global and National, was also likely to impact whether they considered themselves to enjoy elite status, with law firm partners in particular, considered members of an 'ultra elite' (Zuckerman, 1972), separate from other legal and non-legal professionals working within the same firm.

Recognising the likelihood of there being asymmetries of knowledge and power between myself and my participants was, therefore, an important factor to consider when designing my methods for generating data. This meant in contrast to most research where it is the researcher who enjoys the status of subject-matter expert and who, therefore, controls the interview, I needed to be less directive in my approach (Petintseva, Faria and Eski, 2019). This meant while my research questions set the overall agenda for the interview, I adopted a less structured approach to questioning that was designed to promote a more conversational tone; and demonstrate my willingness to let my participants teach me about the phenomenon of AI-enabled legal services (Van Audenhove and Donders, 2019). A successful interview would, therefore, be one in which the participant felt comfortable to disclose information that might not otherwise be public knowledge, and through this help demystify how AI-enabled legal services are experienced by legal professionals.

While my participants could be considered members of an elite, this did not mean they were necessarily experts on the topic of AI-enabled legal services. Unlike elites whose power and status frequently stems from their position within a profession or organisation, expert status is reflective of the knowledge an individual possesses. Hence, experts can be identified through their responsibility for the development, implementation or control of policies and processes (Van Audenhove and Donders, 2019). In the context of my research, expert knowledge of AI-enabled legal services was more likely to reside amongst the *producers* of AI-enabled legal services, than those who *consume* them (Armour, Parnham and Sako, 2022); as these were the individuals responsible for the development, monitoring and evaluation of AI-enabled legal services. These anticipated variations in expertise amongst my participants, therefore, influenced my approach to data generation, as I needed to use methods suitable for both expert and non-expert audiences.

To help prepare myself for the participant interviews, I first held exploratory interviews (Bogner, Littig and Menz, 2009) with my organisational sponsors in Global and National, both of whom possessed expert knowledge of AI-enabled legal services and held *producer* roles within their respective firms. I used these interviews to capture *technical knowledge* about the LegalTech each firm was using (the AI software was demonstrated to me using Zoom's screen sharing functionality), and *process knowledge* about the ways in which AI-

enabled legal services were currently being used by legal professionals at each firm. This helped me better understand the specific context of AI-enabled legal services at each firm. This was also knowledge that I could not assume would be possessed by the participants in *consumer* roles that I might interview. Gaining access to this knowledge in advance of the participant interviews, I anticipated would allow me to better demonstrate my competence, and more easily build rapport and trust with my participants (Van Audenhove and Donders, 2019).

In contrast, my participant interviews required me to tap into the *explanatory knowledge* of my participants, primarily their subjective interpretation of their experience with AI-enabled legal services and the meaning they attached to these experiences. Such knowledge is typically tacit and uncodified, meaning the participant may not be aware that they possess this knowledge. To help participants share this knowledge with me, I tempered the more open interview design that I used with the organisational sponsors, who possessed a high-level of expertise, and instead adopted a semi-structured approach to data generation. Prior to the interviews I asked each participant to complete two reflexive exercises, designed to capture their attitudes towards and experiences of, AI-enabled legal services. These were used as an initial stimulus for discussion in the interview, in addition to which I used a series of follow-up questions where necessary, to ensure that data relating to my research questions would be generated (Bogner, Littig and Menz, 2009).

While all researchers need to be aware of the limitations of interview data (Atkinson and Silverman, 1997; Alvesson, 2011), elite interviews can further exacerbate some of these problems (Petintseva, Faria and Eski, 2019). I considered all my participants to be members of *the interview society* (Atkinson and Silverman, 1997); in addition to this, I was also aware that some of the legal professionals I interviewed were experienced interviewers in their own right, with first-hand experience of using interviews for recruitment purposes and gathering information from witnesses. This meant the participants were aware that the interview afforded them the opportunity to construct responses that would further their own objectives, which may not have aligned with the aims of my own research. Hence, I did not automatically regard my participant's responses as a mirror of reality (or even how that individual perceives reality). Instead, I reflected on the potential meaning of the data and

what claims I could make, using a set of different interpretive possibilities (Alvesson, 2011). Having recognised the elite status of my participants posed additional challenges to using interviews to generate data I utilised several of the metaphors identified by Alvesson (2011) to help me think about how my participants might understand both the interview and my research. These reflections subsequently informed both the design of my interview and the ways in which I interpreted the data generated in the interview.

The privileged position of elites can mean they are protective of information concerning their behaviour and the organisation they work for, and concerned about giving responses that could lead them to be perceived as incompetent (Petintseva, Faria and Eski, 2019). This can lead to responses that seek to portray the participant (and their organisation) in a positive light, for example by providing partial accounts or engaging in post hoc reasoning to explain certain outcomes. While recognising such behaviour cannot be eradicated, I designed the interview to reduce the perceived need for participants to engage in such behaviour. First, I stressed to all participants the limited ways in which the data would be reported and that both they and their organisation would be anonymised. I also used the preamble to the interview (Appendix 1) to acknowledge the different uses of interviews that participants might be familiar with, contrasting the use of interviews by journalists who seek to 'catch out' an interviewee with the more even balance of power that characterises research interviews (Salmons, 2014) and my aim to understand their lived experience. This was designed to reduce the ambiguity that the participant might have been experiencing, help them with their sense-making and reduce the complexity of the interview as a social interaction (Alvesson, 2011). I also recognised that participants were likely to undertake identity work and invoke different identities to guide their responses. Rather than seeing this as problematic, or something that I could control, I instead treated it as a potential object of study (Alvesson, 2011), and asked each participant to reflect on whether they felt their professional identity had changed as a result of their experience of using AI-enabled legal services.

After completing each interview, and prior to conducting a detailed analysis of the data that was generated, I reflected on my experience of the interview, capturing my thoughts about what I believed the participant had been trying to achieve in a memo. For example, did they

try to present themselves in a favourable way, or rely on the use of familiar cultural scripts or discourses to answer my questions? These memo data were something I returned to when later evaluating what degree of confidence I felt in making claims based on the data that had been generated.

Ethical considerations

Ethical approval for my research project was obtained from the University of Sheffield Research Ethics Committee prior to any data being generated. In addition to meeting the requirements of the university ethics policy (University of Sheffield, no date), the research was also designed in accordance with the British Psychological Society's Code of Human Research Ethics (British Psychological Society, 2021a) and their more specific Ethical Guidelines For Internet-Mediated Research (British Psychological Society, 2021b), reflecting my professional membership of this organisation. A summary of the main ethical considerations that impacted the design of my research are detailed below and organised using the British Psychological Society's four research principles.

1. Respect for the autonomy, privacy and dignity of individuals, groups and communities.

To satisfy this principle, researchers are required to develop and follow procedures that protect the rights of participants. To ensure I obtained valid consent from all participants, I provided them with a detailed Information Sheet (Appendix 2), designed to give them the information they would need, in order to determine whether they wanted to participate. This included the overall purpose of the research; who was being invited to participate; the types of data I was seeking, and the methods that would be used to generate data; the time commitment expected from participants and how the data would be stored and used in the future. It was stressed that participation was voluntary, and that candidates could withdraw from the research at a later date, where practical. A Consent Form (Appendix 3) was used to capture the participant's informed consent to take part in the research. In addition to this, participants consent was also recorded at the start of their interview, prior to any data being generated.

As the potential participants were highly educated professionals whose job roles granted them relatively high-levels of autonomy, they were not considered to be vulnerable. In spite of this, I recognised that participation in the research could potentially harm their reputation and career. For example, if a participant shared views deemed detrimental to their law firm or powerful stakeholders within the firm, that later became public knowledge. There was also a risk that the influence of the organisational sponsor could prevent participants from giving truly informed consent, for example, an individual might feel obliged to consent because the invitation to participate in the research was sent by a senior individual within their firm. I was able to largely mitigate these risks through careful research design. By asking interested participants to direct their interest in taking part in the research to me (and not the organisational sponsor) it was possible to maintain the privacy of individuals. Participant data generated during the research was by default anonymised, for example, through removing references to specific processes, roles or entities that might allow an individual to be identified, when combined with other publicly available information. Participants were also given the option to request that their personal details were linked to their comments, in order to allow their voice and contribution to the research to be recognised; however, in practice no participants requested this. A further safeguard to participant confidentiality related to the storage of digital data. Where possible, all data was stored using secure local devices, rather than storing data using networks or systems that I was not in full control of. In those instances where data was generated or transcribed through external parties, I ensured that all data was handled in accordance with General Data Protection Regulation (GDPR) guidelines.

2. Scientific integrity

This principle requires that research should be conducted in ways that ensure a contribution to the development of knowledge, the efficient use of resources and does not devalue the contribution of participants. I sought to satisfy this principle by focusing on a phenomenon of significant interest (AI-enabled legal services) to both academics and practitioners working within this field. I also actively engaged in the processes used by the University of Sheffield to ensure the high-quality of my research, including the Confirmation Review process and regular discussions with my academic supervisors.

The decision to generate data using digital methods meant the physical distance between myself and my participants did reduce my level of control over the conditions in which the research was conducted (BPS, 2021b). I sought to minimise the impact of this by providing participants with guidance on how to maximise the quality of the data that was generated, for example through the use of a high-quality microphone in a setting with low levels of background noise. I also used methods of analysis that did not require the data to have been generated in experimental conditions in which the control of variables is crucial to ensuring validity.

3. Social responsibility

Management research takes place in a context in which there are likely to be differential interests and power imbalances between stakeholder groups. This raises the risk of othering where the consequences of research can be seen to further the goals of one group at the expense of another. This is particularly problematic when this can be seen to perpetuate existing inequalities. Recently I have become increasingly aware that management research has been criticised for developing theories that have influenced the introduction of management practices that prioritise the needs of management and shareholders, over those of employees and the wider community.

To reduce the risk of othering researchers need to be reflexive and transparent about whose perspectives and interests are being foregrounded in their research. I chose to conduct my research at the individual (rather than organisational) level, with the aim of contributing to the 'common good' amongst legal professionals as a whole, rather than focus on the agenda of those responsible for the management of law firms. This meant being reflective about my role as researcher and the possible consequences of my research.

I, therefore, adopted a theoretical stance and identified research questions capable of bridging the theory-practice gap (Van De Ven and Johnson, 2006), meaning I wanted to make both a theoretical contribution to the understanding of AI-enabled legal services and influence their use amongst practitioners. When recruiting participants, I adopted broad inclusion criteria being careful to ensure that I did not exclude the views of different groups within the organisation, so that a diversity of experiences would be represented. I also used

methods to generate data that would be readily accessible to legal professionals in a range of different roles, and which did not limit access to specific groups within National and Global.

I was also explicit in acknowledging the active role I have played in the interpretation of my findings and that my analysis will have been influenced by my personal background and professional experiences within the legal sector. This can be seen through my identified implications for practice, which are designed to make a positive contribution to the experience of those individuals using AI-enabled legal services as part of their professional practice, not just further the goals of the law firms that employ them.

4. Maximising benefit and minimising harm

The likely impact of my research on participants was something I considered throughout the research process. My overall aim was that the research should have a positive impact on the working lives of my participants, while not exposing them to a risk of harm greater than that encountered through their usual pattern of work.

My research topic did not relate to sensitive subject matter, and so was unlikely to cause any emotional or psychological harm to either myself or my participants. The semi-structured nature of my interviews also meant, while the precise nature of the questions varied somewhat between interviews; the purpose of the research, wider themes of questioning, and use of the findings, were all known to participants in advance i.e. no deception of participants was necessary. This meant participants could make an informed choice as to whether they wished to participate. Candidates were also given the opportunity to address any concerns to me prior to signing their consent form and at the start of the interview.

The decision to generate data using digitally mediated methods, meant I did not need to visit any organisational premises or meet participants in person. This ensured COVID-related health risks that could arise from coming into close contact with persons, were avoided. Conducting the interviews electronically using Zoom software, also meant that the interview experience was broadly comparable to the daily experiences that participants had

working with their colleagues and clients during the Pandemic. The use of electronic interviews also gave participants the autonomy to decide where and when the interviews took place. This meant they could minimise the impact that my research would have on both their professional life and personal life, for example by choosing to conduct the interview in a private meeting room at their workplace; or in the event that they were working from home, in a location away from other family members.

4.4 Methods for generation of data

Having considered the above factors I developed a mixed-method approach to generate data that would be suited to exploring the phenomenon of AI-enabled legal services and explicate the causal mechanisms linked to their adoption. The chosen combination of methods is outlined below and was influenced by Wynn and William's (2012) five methodological principles discussed in chapter 3, and guidelines for researching technological phenomena using mixed methods (Venkatesh, Brown and Bala, 2013; Zachariadis, Scott and Barrett, 2013).

While I placed an emphasis on the use of qualitative methods, due to their ability to generate rich and detailed insights (Brönnimann, 2022), I also recognised the descriptive role that quantitative data can play in critical realist research (Zachariadis, Scott and Barrett, 2013). For example, it can be used to quantify the characteristics of a structured entity, and the frequency with which contextual conditions or events are observed within a population. This meant quantitative data would be useful to me when looking for evidence of demi-regularities, but that I would not use it to demonstrate correlations between variables, or present it as evidence of causality, as happens in positivist research.

Data generation with organisational sponsors

Electronic semi-structured Interview

The interviews with my organisational sponsors in Global and National were an opportunity to speak with a subject-matter expert to capture technical knowledge about *what* LegalTech

each firm was using, and process knowledge about *how* it was being used. This data would be important in helping me determine which assumptions of the technology as tool perspective were challenged by the AI-enabled legal services at their firm. In addition to this, I was also interested in context knowledge about the reasons *why* each firm had chosen to adopt AI-enabled legal services. I anticipated that this data would already be known to the sponsors (or their close colleagues), even if it had not been codified in documentary form. To help ensure the sponsors would be able to share accurate data with me, I provided them with a list of questions in advance, so that they could gather relevant data they would need to answer my questions (Appendix 4). These initial questions were then subject to more detailed probing in the interview.

The interview was conducted electronically via Zoom and was semi-structured in nature, and lasted approximately one hour. As experts in the use of AI-enabled legal services at their firm, I gave the sponsors a high-degree of freedom in how to share the requested information with me, so that they would have the opportunity to discuss other aspects of AI-enabled legal services that I may not have anticipated. This meant I used my list of questions as a checklist, rather than a script to drive the interview; where necessary I also probed and asked follow-up questions to the initial responses provided by the sponsors.

Observation of AI-enabled Legal Services

After answering my initial questions, the organisational sponsors allowed me to observe the different AI-enabled legal services in use at the firm. Using Zoom's screen-sharing functionality, I observed the sponsors using the same tools and processes that were being utilised by the legal professionals I would be interviewing. This gave me the opportunity to *listen for the invitational quality of things* (Adams and Thompson, 2016) and generate data about AI-enabled legal services that would not have been possible from an interview alone. More specifically, by observing first-hand the different steps of each process, I was able to better see the different points at which technology and legal professionals intersect. I was also able to observe the affordances of the technology, and see how AI-enabled legal services were framed for legal professionals. This gave me an insight into how the design of the process could impact the behaviour of users, through inviting or discouraging them from taking certain actions. To protect client and firm confidentiality the demonstrations used

dummy data and were not recorded. I was, however, given permission to take notes of what I observed and transcribe these into a written memo for future reference, and used alongside the participant-generated data for analytical purposes.

Data generation with legal professionals

Reflexive Exercises

Prior to participating in an electronic interview, I asked all participants to complete two written exercises, which were designed to help them reflect on their experience of AI-enabled legal services. The exercises were designed as a complement to the electronic interview, which all candidates undertook after they had returned their completed written exercises to me. Prior to the interview I reviewed the content of the written exercises, noting areas where further discussion would be beneficial.

A key aim of the research was to provide participants with an opportunity to share their explanatory knowledge of AI-enabled legal services. However, I recognised that participants can struggle to recall and articulate their tacit knowledge of a phenomenon, when asked direct questions in the context of an interview (Kempster and Parry, 2014). By providing a greater level of structure and time for reflection, narrative exercises completed in advance of an interview can help overcome this limitation. Drawing on Kempster (2006) I developed a reflective exercise (Appendix 5) that asked participants to create a narrative timeline of the experiences that had shaped their attitude towards AI-enabled legal services. To increase the likelihood that the data generated would help me to understand the causal powers and contextual factors that had influenced their decision to adopt AI-enabled legal services, I provided detailed instructions on how to create the timeline and provided an example of a completed timeline.

In addition to the narrative timeline, I developed a further written exercise using Adams and Thompson's (2016) heuristic *Applying the laws of media*. The exercise asked participants to think about the positive and negative ways in which AI-enabled legal services had impacted their work to date, and how they thought things might change in the future. Asking participants to reflect on this using the *four laws of media* (McLuhan and McLuhan, 1988), was intended to generate data that I would be able to analyse, and potentially transform

into quantitative data that could be used for descriptive purposes. This would make it easier for me to identify points of similarity and difference between the experiences of participants, which would help me understand the environmental effects that AI-enabled legal services had on those using them.

Electronic semi-structured Interview

In-depth, qualitative interviews are the primary method used by social science researchers to generate data that describe an individual's experiences and way of understanding the world; and from which analysis of causal relations can take place (Brönnimann, 2022). This made interviews a logical method for generating participant data capable of answering my research questions. The decision to conduct the interviews electronically (rather than face to face), reflected the COVID-19 restrictions discussed earlier.

There are two metaphors I found useful to characterise my understanding and approach to interviewing participants in this research. The first, was to see the interview as a *social encounter* (Holstein and Gubrium, 1995) in which the generated data was co-produced by the participants and myself. Within this encounter, as interviewer I adopted the role of *gardener*, planting the seeds of the interaction through the questions I asked; and cultivating and shaping the participant's response through my behaviour and the use of follow-up questions (Salmons, 2014).

In adopting this approach, I openly acknowledged the work that is undertaken by both interviewers and interviewees to generate data, which in turn requires data to be critically evaluated, before claims are built upon it (Atkinson and Silverman, 1997). In seeking to avoid the charge that I was uncritical of my participants' accounts, I designed the interview with critical realist underpinnings (Brönnimann, 2022). This meant while I was interested to learn about my participants' experiences and reflections of AI-enabled legal services, this was not the only goal of the interview. Instead, these accounts also served as a means for me to identify the social conditions under which my participants' experiences took place, and how these conditions had led to the adoption of AI-enabled legal services.

The semi-structured nature of the interview was designed to give participants a degree of flexibility in how they discussed the topic of AI-enabled legal services. However, because not all participants could be considered experts on the topic, I retained a greater degree of control than in the organisational sponsor interviews. Hence, to maximise the likelihood that the interviews would generate the data necessary to answer my research questions, each interview followed a similar sequence (Appendix 1).

At the start of the interview, I asked participants to provide biographical data relating to their current role and their experience of using AI-enabled legal services. I then progressed to an exploration of AI-enabled legal services within the context of the participant's firm. To generate data that went beyond mere description of empirical experiences, several of my questions focused on getting participants to explain *how* and *why* their experiences had taken place (Wynn and Williams 2012). For example, '*How did AI become part of your professional practice?*'

Through reference to the written exercises of the participant, we then discussed how their attitude towards AI-enabled legal services had developed, and how their work was being impacted. The ability to discuss the written exercises during the electronic interview gave participants the opportunity to elaborate upon and clarify their experience of AI-related legal services. This generated further data that I could analyse to identify the different entities, causal powers, and contextual factors that led to their decision to adopt AI-enabled legal services at work. Having reviewed the exercises in advance, the discussion also allowed me to share my tentative interpretations of what they had written and the wider themes that were emerging in my research. This gave me an opportunity to see whether my ideas resonated with their own experience, and allowed them to challenge what they felt was inaccurate, which helped refine my understanding of the phenomenon.

To explore how my participants understood their relationship to artificial intelligence at work I asked them to reflect on whether one or more of Ihde's, (1990) four types of human-technology relations captured their experience of using AI-enabled legal services. This question was designed to create the opportunity to see both the human and technology at once (Adams and Thompson, 2016), through understanding how AI-enabled legal services mediated how legal professionals interacted with the world.

The remainder of the interview was used to discuss trust and professional identity, and how these concepts related to the participants' experience of AI-enabled legal services. My aim was to understand whether the participants trusted AI-enabled legal services and, if they did, how their *initial trust* and *experience-based trust* had developed. In relation to professional identity, I was interested to know whether the participants felt their identity had shifted since using AI-enabled legal services and, if it had, what was different as a result.

I closed the interview by asking each participant whether they wished to revisit any of their comments, to give them an opportunity to elaborate or correct any comments they had made earlier in the interview. Reflecting the explorative nature of the research, I also asked them whether they thought there were any important aspects of AI-enabled legal services that we had not touched on. Finally, having completed their interview and learned more about the research, I asked them if they were aware of any individuals at the firm that would be knowledgeable of the topic, and who I should approach to interview.

5. ANALYTICAL STRATEGY

Guided by the general maxims of causal case study research, my analysis was undertaken in two distinct stages, each of which employed a different method of analysis designed to answer different questions. The first stage utilised a critical realist informed thematic analysis (Wiltshire and Ronkainen, 2021; Fryer, 2022) to answer the question, *How are AI-enabled legal services understood by the legal professionals that use them?* Building on this, the second stage of analysis employed theory-building process-tracing (Beach and Pedersen, 2016) to answer the question, *What explained the use of AI-enabled legal services amongst these legal professionals?* A within-case analysis of both a typical *producer* case, and a typical *consumer* case produced two empirical narratives (one for each case) from which it was hypothesised what causal powers and mechanisms could best explain the use of AI-enabled legal services. A cross-case analysis was then used to determine the extent to which the hypothesised causal pathways were supported by the empirical data of contextually similar cases.

5.1 General principles for rigorous critical realist case study analysis

Having made the decision to adopt a case study approach, it was important to conduct it in a rigorous way that reflected what is considered best-practice (Piekkari, Welch and Paavilainen, 2009) For case study research to be considered rigorous it needs to satisfy several widely accepted markers of rigour, such as reliability, inference quality, construct quality, and generalization (Avenier and Thomas, 2015). However, it is also important to recognise that the precise meaning of these terms, and as a corollary the analytical methods for evaluating them, are dependent on the theoretical position of the researcher. It is, therefore, necessary to explain the use of these terms in the context of critical realist case study research (Table 14). There also needs to be transparency about how the data was generated and the specific steps that were taken during the process of analysis. These general principles informed the strategy for analysis (see section 5.3).

Table 14: Markers of rigour in critical realist case studies

Reliability	In critical realist research reliability describes the transparency and clarity of the cognitive path taken by the researcher when moving from the empirical material to the findings of the research. For their research to be considered reliable, critical realist researchers need to explain in detail how they have formulated their research questions, generated, transcribed and interpreted their data, and reported their results.
Quality of inferences	Inference quality is evaluated through reference to how plausible are the explanations given for any differences or regularities observed in the data. Such explanations will comprise configurational patterns of events, structures and mechanisms linked to the phenomenon. It is also important that the researcher explains the process through which such patterns are generated.
Quality of constructs	Construct quality depends upon the explanatory power of the model it is a part of. In critical realist research constructs comprise the structures and wider context that through their interaction activate the observed pattern of events. Evaluating construct quality involves determining how well the model accounts for what has been observed empirically. This means construct quality is also dependent upon the quantity and precision of the data on which the model is based. High-quality constructs allow for the development of explanatory models that are both parsimonious and better able to explain the observed phenomenon than alternative theoretical models.
Generalization	Generalization concerns the degree of abstraction of the generated explanatory model. Critical realist research does not generalise findings through statistical inference, or aspire to predict events elsewhere. Instead, generalising means identifying common structures and mechanisms theorised to exist elsewhere. Critical realists utilise knowledge gained through identifying structures and causal mechanisms in one context, to explain similar events

Table 14: Markers of rigour in critical realist case studies (continued)

	elsewhere (<i>demi-regularities</i>). Comparative case methods can be used to identify additional contextual and structural factors that impact the workings of causal mechanisms, with the aim of deepening understanding, rather than seeking to falsify or enhance the validity of the mechanism.
--	--

Adapted from Avenier and Thomas (2015)

5.2 Theorising from case study data

Contextualised explanation

In addition to the general principles above it was also important to understand the way in which critical realist case studies can be used to theorise, given that the primary focus of the research was to provide a theoretically plausible answer to the question, *‘What explained the use of AI-enabled legal services amongst these legal professionals?’*

Critical realist case studies can provide this type of theoretical explanation by generating a contextualised causal explanation of empirical events linked to a phenomenon (Welch *et al.*, 2011, 2022). In this approach to theorising, a strong emphasis is placed on both context and causality, with cases used as sites to trace the workings of the mechanisms that produce a specified outcome, in this case the use of AI-enabled legal services. This requires the distinction that is made between the phenomenon and its context in positivist case study research, to be dissolved (Welch *et al.*, 2021). Instead, context is seen as more than just a descriptive container within which the phenomenon sits, rather it is treated as having explanatory power. In my own research, this meant considering how the specific context of a law firm and the wider legal profession might be contributing to AI use.

The distinctiveness of the contextualised explanation approach to theorising from case studies, is clearer when compared to methods of theorising employed in case study research that is underpinned by alternative theoretical standpoints (Table 15).

Table 15: Four methods of theorising from case studies

<i>Dimension</i>	<i>Inductive theory building</i>	<i>Natural experiment</i>	<i>Interpretive sensemaking</i>	<i>Contextualised explanation</i>
Philosophical orientation	Positivist (empiricist)	Positivist (falsificationist)	Interpretive/ constructionist	Critical realist
Nature of research process	Objective search for generalities	Objective search for causes	Subjective search for meaning	Subjective search for causes
Case study outcome	Explanation in the form of testable propositions	Explanation in the form of cause–effect linkages	Understanding of actors’ subjective experiences	Explanation in the form of causal mechanisms
Strength of case study	Induction	Internal validity	Thick description	Causes-of-effects explanations
Attitude to generalisation	Generalisation to population	Generalisation to theory (analytic generalisation)	“Particularisation” not generalisation	Contingent and limited generalisations
Nature of causality	Regularity model: proposing associations between events (weak form of causality)	Specifying cause–effect relationships (strong form of causality)	Too simplistic and deterministic a concept	Specifying causal mechanisms and the contextual conditions under which they work (strong form of causality)
Role of context	Contextual description a first step only	Causal relationships are isolated from the context of the case	Contextual description necessary for understanding	Context integrated into explanation
Main advocate	Eisenhardt	Yin	Stake	Ragin/Bhaskar

Source: Welch *et al.*, 2011 (p.745)

The contextualised explanatory accounts generated using a critical realist case study should encompass both human intentionality, as articulated by the case participant(s), and their position in the social structure. Hence while it is meaningful to talk of human agency, a contextualised explanation needs to go beyond the understanding and interpretation of individuals, to also account for wider factors, which may not form part of their account (Welch *et al.*, 2011).

Drawing on critical realism, the approach also challenges the idea of purely inductive or deductive theory development. Instead, the process of explanation frequently starts with an anomaly, an observed outcome that is different from what was anticipated. Anomalies can arise for various reasons, for example, the existence of a new causal factor, a lack of understanding about the case context, or limitations about the precise workings of a causal mechanism (Welch *et al.*, 2011). This mismatch between empirical observations and

existing theory requires the phenomenon of interest to be reconceptualised in a way that accounts for the anomaly. The generation of causal explanations utilises retroductive thinking - working backwards from the observed outcome, to identify the causal process that explains how the outcome could have arisen from the context. More specifically, in case study research this type of research utilises a process frequently referred to as 'process tracing' (George and Bennett, 2005), to develop an explanatory narrative that can draw upon existing theories and known patterns, to identify the links between the different elements in the causal chain.

When theorising in this way, there is also a recognition that combinations of conditions may be required in order for an outcome to occur i.e. the links in the causal chain are not acting in isolation from one another. Cases and the elements that make them up, must, therefore, be considered holistically. There may also be a role for equifinality i.e. the same outcome emerging from different combinations of conditions. Thus, context is treated as central to the explanatory account, rather than being treated as 'idiosyncratic detail' as it would be in more positivist case approaches, such as those of Eisenhardt and Graebner (2007). The approach, therefore, sits in clear contrast to correlational models of causation that seek to demonstrate explanation by calculating the strength and direction of the relationship between independent and dependent variables, but which can lack explanatory power for the processes involved.

In summary, contextualised explanation helps overcome an historic limitation of case study research that regarded the method as suitable for inductive theory-building (but only at the expense of context). This is achieved through producing causal explanations that are more contingent and which should only be generalised within similar contexts. Hence, despite its more limited scope, the benefits offered by the contextualised approach meant it was well-suited to answering the question, *'What explained the use of AI-enabled legal services amongst these legal professionals?'*

Different types of contextualised explanation

Within the broad category of contextualised explanations, a variety of different approaches exists. Of the different approaches highlighted by Welch *et al.* (2022) *historical research*

was not well suited to the investigation of a contemporary phenomenon where there was limited historic data available to draw upon. The *extended case study method*, which draws on the ethnographic tradition, was not feasible to implement at the time I was generating my data, given the restrictions placed on field research during the COVID pandemic. The remaining two approaches – *process research* and *configurational logic* - were both plausible candidates for developing an explanation for the use of AI-enabled legal services. However, the analytical techniques that underpin configurational logic perform best with a somewhat heterogeneous case population, that includes examples where the outcome of interest has and has not taken place. The homogeneity of my data meant that my case population only comprised instances in which AI-enabled legal services had been accepted. This meant I was unable to compare examples of acceptance and non-acceptance, thus limiting the benefit of a configurational logic approach. I, therefore, developed an approach that drew on the logic of *process research*. This approach employs an exploratory style of theorising and generates causal explanations through narratives and models, that capture the process and its underlying mechanisms, within their wider context (Welch *et al.*, 2022).

5.3 Overview of the data analysis process

To pursue my phenomenon-based research strategy, which was focused on distinguishing, exploring and theorizing aspects of AI-enabled legal services, two different types of analysis were identified as relevant. First Wiltshire and Ronkainen's (2021) method of thematic analysis, which is underpinned by the critical realist concept of ontological depth, was used to distinguish and explore the phenomenon of AI-enabled legal services, identifying patterns of similarity and difference across the population of cases. The findings of this analysis were then reorganised to generate a technological frame of AI-enabled legal services that reflected the participants' understanding of the technology. The analysis also provided the opportunity to identify contextual information relating to each case, which proved useful when seeking to understand how differences between the cases had emerged.

This method of thematic analysis generated three different types of themes – experiential, inferential and dispositional – which built on one another, becoming increasingly abstract.

This allows the analysis to move from an initial focus on description of the empirical material to a more interpretive focus, which was used to conceptualise my understanding of AI-enabled legal services. While not the sole focus of this stage of the analysis, Wiltshire and Ronkainen's (2021) method also provided the opportunity to undertake some initial theorising as to what might have caused the events identified in the thematic analysis. This type of causal theorising generated what Beach and Pedersen (2016) call a *minimalist* understanding of causality. This meant while the link between a cause and effect was acknowledged, the theorised mechanism linking the two was not unpacked in detail, which limited the strength of the causal claims that this stage of the analysis could make. While this can be seen as a limitation of the method, it is not realistic to undertake an in-depth causal analysis of every research finding (Beach and Pedersen, 2016). I therefore sought to unpack the causal mechanisms that had the potential to explain the key findings of the thematic analysis, using, process tracing (Beach and Pedersen, 2016), as a further method of analysis.

This involved using the coded and contextualised data from the thematic analysis as the raw material for process tracing. This stage of the analysis focused on providing a detailed and contextualised causal explanation of the use of AI-enabled legal services amongst legal professionals. This more comprehensive process of causal analysis, which used both within-case and cross-case analysis, was designed to generate a deeper *systems* understanding (Beach and Pedersen, 2016) as to why the legal professionals in my research had accepted the use of AI-enabled legal services, and made them a regular part of their professional practice. A detailed discussion of the steps involved in the two different methods of analysis is provided below.

5.4 Process of thematic analysis

Related research question: How are AI-enabled legal services understood by the legal professionals that use them?

This research question was chosen to develop my understanding of the ways in which legal professionals experienced AI-enabled legal services, for example their knowledge, assumptions and expectations of the technology, which could later be used to develop a technological frame for AI-enabled legal services (Orlikowski and Gash, 1994). At the same time, I also recognised that their interpretations were shaped by both material reality and the wider context in which the legal professionals were working. This led me to identify thematic analysis as a suitable methodological approach that would allow me to describe and interpret my data in a way that could answer this question.

Thematic analysis is regarded as a foundational method of qualitative analysis, as it requires researchers to develop a set of skills that underpin qualitative research more generally (Braun and Clarke, 2006). I determined that a thematic analysis was a pragmatic way for me to conduct an initial analysis of my data; the aim of which was to organise the data in a parsimonious way; provide a thick description of AI-enabled legal services from the perspective of legal professionals; while also affording me the opportunity to interpret the data for meaning using theoretical insights from the domains of information systems, psychology, and management. Ultimately, the chosen thematic analysis would identify themes (including the highlighting of similarities and differences) across the data set as a whole, rather than looking for themes that occurred within individual cases. Thereby offering an insight into how legal professionals experience AI-enabled legal services.

While thematic analysis was not the only analytical method that could be used to achieve these goals, I felt it offered several advantages when compared to potential alternatives. Given the embryonic nature of research on AI-enabled legal services, thematic analysis was attractive as it offered greater freedom when analysing my data than more deductive methods, such as content analysis, which typically require researchers to establish a priori rules for coding data prior to the start of the analysis. Grounded theory was another option I considered, as it too offered a more data-driven approach to analysis, it also has the benefit of having been adapted to different philosophical standpoints, including critical realism (Hoddy, 2019; Looker, Vickers and Kington, 2021). However, the grounded theory approach is regarded as a challenging method for novice qualitative researchers to utilise. Aligning the grounded theory method to the institutional requirements of PhD research,

which typically requires researchers to review the extant literature prior to submitting a formal research proposal, is also problematic given grounded theory should be data rather than theory driven (Thurlow, 2020).

Thematic analysis has the advantage of being compatible with a range of ontological and epistemological positions. This spectrum of approaches to thematic analysis ranges from those with positivist assumptions e.g. coding reliability thematic analysis (Boyatzis, 1998) to those that reflect constructionism e.g. reflexive thematic analysis (Braun and Clark, 2006). This provides researchers with a degree of flexibility as to how the method is implemented; provided the assumptions that underpin the chosen approach to thematic analysis are in alignment with the research questions they are seeking to answer and the researcher's theoretical stance (Braun and Clarke, 2006). However, it can mean that different approaches have little in common, once they are looked at in detail (Fryer, 2022).

In the context of critical realist research, it is only recently that specific methods of thematic analysis have been proposed (Wiltshire and Ronkainen, 2021; Fryer, 2022). These approaches have sought to overcome tensions that have arisen between the positivist and constructionist approaches to thematic analysis. This has seen the development of approaches that can accommodate the critical realist assumption of ontological depth, through the use of methods that allow for both 'surface-level' (semantic) description and 'deep-level' (latent) interpretation and reflection, both of which are important when seeking to develop knowledge about a phenomenon (Wiltshire and Ronkainen, 2021). This means that while both Wiltshire and Ronkainen's (2021) and Fryer's (2022) methods have points of commonality with other approaches to thematic analysis e.g. reflexive thematic analysis (Braun and Clark, 2006), they develop themes in distinctive ways that reflect their critical realist underpinnings.

While both critical realist approaches have their individual merits, Fryer (2022) indicates his narrower approach, which focuses explicitly on generating themes that provide causal explanations, is better suited to explanatory research projects within well-documented fields. In contrast, Wiltshire and Ronkainen's (2021) approach suggests the use of three types of theme – *experiential*, *inferential* and *dispositional* – that can be used for descriptive purposes as well as theorizing. These themes reflect the three levels of Bhaskar's stratified

ontology. Experiential themes are used to describe participants’ subjective understandings and experiences, for example, their intentions, feelings and beliefs, as they appear in the data. In contrast, inferential themes describe the interpretations a researcher or a participant may make about aspects of the social world, which have not been observed empirically, these take the form of inferences and conceptual redescriptions. Finally, dispositional themes are attempts to theorise about the causal powers and mechanisms that must exist for the phenomenon of interest to manifest.

The flexibility of Wiltshire and Ronkainen’s (2021) approach was attractive to me given the exploratory nature of my research about AI-enabled legal services and the fact that the research question the thematic analysis was being used to answer focused on the experiences of legal professionals using AI-enabled legal services, rather than what might be the underlying causes of the phenomenon. I, therefore, adopted Wiltshire and Ronkainen’s (2021) method for the thematic analysis, but integrated Fryer’s processes for standardization and consolidation when developing themes. Prior to beginning this analysis, I also utilised some of the broader principles proposed by Fryer (2022), that reflect wider practices associated with other variants of thematic analysis (Braun and Clark, 2006). This decision can be seen to answer Fryer’s hope, *“that the two models are read together, and perhaps even improved together in conversation”* (p.382). A summary of this integrated approach is detailed below (Table 16), prior to a more detailed overview of each step in the process.

Table 16: Critical realist approach to thematic analysis

General preparation (Fryer, 2022)
<p>Step 1 - Data preparation and familiarization</p> <p>This stage comprised data transcription and data management using NVivo; familiarising myself with the overall data corpus to identify the data within it that would be used for further analysis; recording of initial thoughts and questions.</p>
<p>Step 2 - Refinement of research questions</p> <p>This stage involved the identification of notable outcomes and anomalies in the data that could be used to refine my proposed research questions.</p>

Table 16: Critical realist approach to thematic analysis (continued)

Method for thematic analysis (Wiltshire and Ronkainen, 2021)
Step 3 - Development and evaluation of experiential themes This stage of analysis focused on observed experiences and events, with the goal of developing inductive themes, which described participants' understanding of AI-enabled legal services in lay-terms.
Step 4 - Development and evaluation of inferential themes This stage of analysis focused on identifying unobserved but occurring experiences and events, which could be inferred from the experiential themes. The resulting inferential themes described my interpretation of what was happening in more abstract terms.
Step 5 - Development of dispositional themes This stage of analysis focused on the use of retroductive methods to identify the causal powers and mechanisms that could most plausibly explain the events described in the inferential themes.

Adapted from Wiltshire and Ronkainen (2021); Fryer (2022)

Data preparation and familiarization

Interview data preparation

Prior to undertaking a formal analysis of the data generated in this project I sought initially to familiarise myself with the data and get a general sense of what was being said by the participants. This began through the process of transcription, which is often regarded as the first step in qualitative data analysis (Kowal and O'Connell, 2014). The process of transcription is often glossed over in management research, being treated as an atheoretical and mechanical task of writing up (Point and Baruch, 2023). In reality, transcription is a subjective process during which the researcher uses their judgement to determine what data is included (and omitted) or deemed in need of correction. This is, therefore, not a neutral process, as it has a direct impact on how participants are portrayed and what knowledge and information is foregrounded for further analysis (McMullin, 2023).

Transcription should, therefore, be treated as a reflexive process in which rigour is achieved (and error and bias minimised) through producing an accurate and truthful representation

of the original data, using a process that is aligned with the researcher's overall approach and their research questions (Point and Baruch, 2023).

The type of transcription I undertook reflected my goal of producing written transcripts that acknowledged the active roles of both me and the research participants in generating data. I, therefore, transcribed all the oral questions, prompts and interjections made by myself, in addition to the content generated by the participants. It was also important that the finalised transcripts reflected the data that I used during my process of analysis, as this would promote transparency and enhance the reliability of my findings. My chosen analytic methods of thematic analysis and process tracing both focused on interpreting the meaning of *what* was being said. This meant I chose to transcribe a thin level of description, what Bucholtz (2000) calls 'intelligent verbatim', which would enable me to connect with and develop an understanding of my participants' experiences. This type of transcription omits the extensive detail necessary for forms of analysis (e.g. conversation analysis) which focus on *how* things are said, by utilising the prosodic, paralinguistic and extra-linguistic aspects of speech.

I managed the process of transcription with the assistance of Nvivo Transcription software, which utilises AI to automatically transcribe digital voice recordings. The interviews were initially recorded using Zoom meeting software; recordings were saved locally in MP4 format on my personal computer, rather than hosted in the 'cloud' by Zoom in order to comply with UK data protection. The MP4 files were then uploaded to Nvivo Transcription, which transcribed the audio to text, for me to review in more detail. After completing my review, copies of both the MP4 file and the accompanying transcript (in .txt format) were uploaded to Nvivo 1.7, where the recording and transcript were synchronised to aid future analysis.

I was initially attracted to using Nvivo Transcription as I had a hunch that the experience of using an AI-enabled service, similar to the phenomenon I was studying, would sensitise me to the working lives of my participants. This would help me anticipate and empathise with the experiences they might discuss in the interview, as at that point in my research I had not had first-hand exposure to AI-enabled services in a work context. On reflection I consider this decision to have been justified, given the experiences I had through my use of AI

transcription. Using an AI-service allowed me to ask myself a similar set of questions to those I planned to use when interviewing participants about their experiences of using AI. I was able to experience first-hand the efficiency gains associated with an AI-augmented process, which allowed me to produce a 'rough-cut' transcription on the day of the interview. The ability to read through a transcription shortly after the completion of the interview was helpful when seeking to record my initial reflections about the interview in a memo (di Gregorio, 2021). I also experienced several of the frustrations associated with AI software in the early stages of development. Every transcription contained numerous inaccuracies e.g. the incorrect use of proper nouns and colloquial language that was not included in the AI's dictionary. After each interview I updated the dictionary with any words that had not been recognised in the previous transcription; this meant the words would be available for future transcriptions. This allowed me to see how the AI was able to 'learn' and improve its performance over time. Other issues arose in relation to participants who spoke with a pronounced regional accent e.g. Glaswegian and when participants did not have access to a high-quality microphone. These were issues that it was not possible for the AI to overcome, thus highlighting the ongoing limitations of the software. This meant I was never able to fully trust the software to perform the transcription task unsupervised meaning the process of transcription can be seen as an augmentation, rather than an automation, of the transcription process. In a comparison of manual transcription and augmented transcription I found that transcription time was reduced by approximately 25% when using the software; significantly less than might have been anticipated given recent research (di Gregorio, 2021; Wright, 2023).

In spite of the insights I gained from my approach to transcription, I recognise the relative novelty of the approach and that similar methods have been subject to critique. Rapid developments in the efficacy of AI transcription have meant the process I employed has only recently become a viable option for researchers. This provides a partial explanation for its limited use to date in management research (Point and Baruch, 2023). However, it is argued that AI transcription is now able to provide 'good enough' first drafts of transcripts to be of genuine use to researchers (Bokhove and Downey, 2018, cited in McMullin, 2023) and can significantly reduce the amount of time spent on transcription (Richardson, Godfrey and Walklate, 2021; Wright, 2023).

These developments also mean that earlier criticisms of auto-transcription in the extant literature should be reconsidered in light of the technology currently available to researchers (McMullin, 2023). A long-standing view amongst qualitative researchers is that manual transcription is an essential step for researchers seeking to immerse themselves in their research, as it allows them to get closer to the original data, prior to its simplification (to a greater or lesser extent) during transcription. While immersion can be compromised by transcription processes that outsource the task away from the researcher to other humans or a technological solution, it is not clear that this criticism necessarily applies to AI transcription. Rather than automating the process of transcription, AI can be better understood as augmenting the work of researchers, as the process still requires significant human involvement, which helps guard against the risk that the tools start driving the research (Wright, 2023). In my research I reviewed every draft transcript Nvivo Transcription produced by reading the transcript while simultaneously accessing the original recording of the interview. This allowed me to ensure accuracy, resolve ambiguities and enhance readability. This was an active process that required me to make decisions in relation to the use of punctuation, whether to indicate non-verbal content and how to treat participant hesitations and repetitions. It also necessitated extensive, direct interaction with the original recorded material. Indeed, this direct interaction continued after the process of transcription was finished as I continued to be able to access the original recordings when undertaking my subsequent analysis (having synchronised the transcript and recording in Nvivo, where I undertook my analysis), something that is much harder to achieve when using traditional transcription methods.

Using AI to augment human transcription can also help reduce the number of errors that arise in manual transcription, even amongst professional transcribers (Poland, 1995). These human errors reflect the arduous nature of the task and the limited capacity of humans to concentrate for long periods. While AI transcription was certainly not error-free in my experience, my ability to review the initial draft with a 'fresh mind' helped me identify the errors with relative ease. The performance benefits of such 'human in the loop' processes, that combine humans and AI, has been demonstrated in several fields, such as medicine (Patel *et al.*, 2019) suggesting an augmented approach to transcription has the potential to offer significant benefits in terms of both accuracy and efficiency to researchers.

It is also important to consider how researcher time saved on transcription is used. While my PhD research afforded me significant time for data analysis, this is not the case for all researchers, such as those employing rapid research methods (Richardson, Godfrey and Walklate, 2021). Where time is at a premium, the ability to reallocate time to allow a more detailed analysis of the transcribed data may enhance the overall quality of research (Wright, 2023).

Having completed the process of transcription and uploaded the recordings and transcript to Nvivo I was able to further familiarise myself with the data. I achieved this by reading each interview transcript several times to get a general sense of 'what was going on', capturing my thoughts in a written 'thoughts and questions' memo within Nvivo. Although a formal process of analysis was not used at this stage, familiarizing myself with the data helped me to home in on what was surprising and interesting. This process of reflection subsequently helped me to refine my research questions for the project.

Reflexive exercise data preparation

Having prepared the interview data, I also reviewed the data that participants had generated through the two reflective exercises they had been asked to complete prior to their interview. The first exercise involved constructing a narrative timeline of key moments relating to the use of AI in the workplace; the second exercise required participants to capture their thoughts about how AI affected their work. In addition to providing participants with an opportunity to reflect in advance of their interview, I also used the data generated through these exercises to inform and evaluate the findings of the thematic analysis generated from the interview data.

Each narrative timeline experience was given a descriptive code (Saldana, 2016), which summarised the experience in either one word or a short phrase. The data were then recorded in a spreadsheet. This allowed me to categorise all the coded experiences reported by the participants and record the frequency with which each type of experience was reported across the population as a whole. Where the data allowed, each experience was also coded for magnitude (Saldana, 2016) in terms of its direction (positive experience; negative experience) and its intensity (high; medium; low). The finalised spreadsheet,

therefore, provided a summary of responses for the overall population, and a detailed overview of the experiences of each participant.

Participants recorded their reflections on how AI affected their work in different categories, reflecting McLuhan and McLuhan's (1988) four laws of media (enhances; obsolesces; retrieves and reverses into). Each participant's data were given a descriptive code (Saldana, 2016), and recorded in a spreadsheet. The coded data were categorised according to which law of media they related to; the frequency of each code was also calculated. As with the narrative timeline, the finalised spreadsheet, summarised the responses of individual participants and the population as a whole.

Taken together these two exercises gave me an important insight into different aspects of the social structure (causal factors) and how they might relate to one another. The narrative timeline was used to highlight participant perceptions as to how different individuals, groups, practices and discourses impacted their attitude towards AI-enabled legal services. The data relating to the four laws of media focused on participant perceptions of AI-enabled legal services as a technological process, thus helping make this visible as another important element within the social structure.

Refinement of research questions

The universal use of AI-enabled legal services amongst the participants in my research was an unexpected finding given that AI is typically used in a small proportion of client work (The Law Society, 2019) and lawyers have been highlighted as more likely to resist technological change than other professionals (Bell, Rogers and Legg, 2019). Anticipating that my data would comprise users and non-users of AI, I originally planned to analyse the data with a view to identifying the causal factors and mechanisms leading to use and non-use. In light of my data, I revised my research question and instead focused exclusively on the question, *What explained the use of AI-enabled legal services amongst these legal professionals?*

I was also interested that many of the legal professionals who were using AI-enabled legal services, struggled to define what AI meant to them in the context of their work. While this was not entirely unanticipated, given that AI remains a contested term, it was interesting

that this did not stop the lawyers talking about the positive attitudes they had developed towards AI-enabled legal services and the experiences that had underpinned this. This confirmed to me the merit in retaining the research question, *How are AI-enabled services understood by the legal professionals that use them?* in its original form.

It was also apparent to me that the level of technological understanding of AI varied significantly amongst the participants, with those working in multi-disciplinary teams generally more knowledgeable than those working in traditional fee-earner roles. This bifurcation reminded me of the distinction made between *producers* and *consumers* of AI-enabled legal services (Armour, Parnham and Sako, 2022); concepts which had emerged following my initial literature review. This led me to consider whether even though similar outcomes were observed amongst all the legal professionals (i.e. the use of AI-enabled legal services and the development of positive feelings towards the use of AI technology), the causal factors and mechanisms that underpinned these outcomes might differ between the two groups. This created two subsidiary questions that the analysis would seek to answer, *What explained the use of AI-enabled legal services amongst the producers of legal services?* and *What explained the use of AI-enabled legal services amongst the consumers of legal services?*

The process of familiarization also confirmed to me that I had been correct to identify trust in technology (and how trust develops) as an important topic for the research to consider. The way in which participants described the role of trust, suggested that it might be acting as a causal mechanism in relation to the decision to use AI-enabled legal services. Hence, in order to understand the role of trust in relation to user acceptance, I would need to unpack it as part of the wider process tracing analysis, rather than treating it as a discrete topic, as originally planned. Hence the research question, *'How do legal professionals develop trust in AI-enabled legal services?* was considered as part of the over-arching question, *'What explained the use of AI-enabled legal services amongst these legal professionals?'*

Development and evaluation of experiential themes

Experiential themes seek to capture the subjective perceptions and experiences of participants. In the context of critical realist research, experiential themes can be seen to

relate to the *empirical domain*, the realm of experiences and events that can be observed. The themes are data-driven and descriptive in nature, as they are produced inductively through identifying patterns in the data, which relate to the participants' intentions, feelings, and beliefs.

In my research the first reading of each interview transcript focused on identifying contextual information about the participant and identifying the participant's experiences. Attribute coding (Saldana, 2016) was used to log the contextual data for each candidate, this was then recorded in a spreadsheet for use during later stages of analysis, but it did not form part of the data used to generate the experiential themes. For example, once an experiential theme was developed, I would then consult the contextual data of those participants whose data contributed to the theme, to see whether there were identifiable patterns amongst the participant data, e.g. did all the candidates whose data contributed to the theme work for a specific law firm or perform a specific job role.

These contextual data included information relating to the participant's gender, organization, job role, level of experience and the length of time they had been using AI-enabled legal services. Typically, this data was generated by the candidate at the start of the interview, in response to being asked to summarise their current role and experience of using artificial intelligence. Where the interview data was incomplete, I sought to identify the missing contextual data in other data sources. For example, the reflective exercises completed by each candidate in advance of their interview typically indicated when they started using AI at work. It was also possible to cross-reference job roles and work histories through publicly available information on the websites of Global LLP, National LLP and LinkedIn.

The experiential themes were generated using the method specified by Wiltshire and Ronkainen (2021). Alongside identifying the contextual data relating to each participant, I coded my tentative ideas about the experiences of each participant using a descriptive code, a word or short phrase, that closely captured what had been said by the participant (Saldana, 2016), while not necessarily quoting them directly (as would be the case when using an *in vivo* code). The aim of coding in this way was to try to capture what the event or experience meant to the participant, and to keep my interpretation at this stage to a

minimum. This would also help promote reliability in the findings of the research by ensuring that subsequent stages of analysis were grounded in the original empirical data.

The now coded interviews were then reviewed, with my focus shifting to the generation of nascent experiential themes that described each participant experience. Each distinct experience that I had coded was recorded in a master spreadsheet as a nascent experiential theme and written using the same format, which adopted the sentence structure advocated by Wiltshire and Ronkainen (2021). This meant each experiential theme began with the identical stem, *'The participant expresses that...'*. Otherwise, the language I used to capture the theme sought to reflect the language of the participants, rather than more abstract language that I might be familiar with from the theoretical literature.

An example of how the originally transcribed data was converted to an experiential theme is provided below.

"So, it is definitely a very positive outlook on technology and AI because I think it just smooths the processes. It helps the clients, it helps the firm and it helps you be in communication with people and it knows what you're doing and where you are in a case."
[Participant 7]

This data was descriptively coded as:

"The participant expresses positive feelings towards the impact of AI on their work."

After reviewing the first transcript, further transcripts were then reviewed one at a time following a similar approach. Where new nascent themes were identified, these were added to the master spreadsheet. In instances where a theme already identified by a previous participant was recorded in subsequent participant accounts each instance of the theme was recorded. When similar experiences were phrased by participants in slightly different ways, I utilised Fryer's (2022) process for standardization to avoid theme proliferation. This would involve adjusting the wording of the original nascent theme to better reflect the experiences of all participants.

Once all the interviews had been reviewed, the finalised list of nascent experiential themes was further analysed to determine the prevalence and strength of each theme across the data set as a whole. Themes that occurred frequently, could be seen as evidence of ‘demi-regularities’ within the specific context of legal professionals using AI-enabled legal services. Alongside the experiential themes, the contextual data was also used to identify patterns within the data set, e.g. differences in response according to job role. Upon completion of this analysis, each nascent experiential theme was rephrased to produce a finalised theme that reflected its existence within the dataset as a whole, rather than with reference to a specific participant. Each finalised experiential theme used the suggested stem, “The data show that [some/many/most] participants in this study [strongly/weakly]...” (Wiltshire and Ronkainen, 2021). For example, the nascent theme,

“The participant expresses positive feelings towards the impact of AI on their work.”

Was identified in the accounts of 20 different participants and in total was recorded in 30 different instances. I, therefore, reworded the theme to reflect its relevance to a majority (most) of the participants.

“The data show that most legal professionals in this study feel positive about the impact of AI-enabled legal services.”

However, given the number of references to the theme was typically just once or twice per participant, and the recorded data did not allow me to accurately determine the participant’s strength of feeling about their experience, I did not feel able to indicate what degree of positivity was experienced towards AI-enabled legal services.

In other instances, the description ‘many’ was used where over a third (but fewer than half) of participants identified a theme. ‘Some’ was used for themes that were identified by less than a third of all participants, but which were deemed relevant to the overall analysis due to the higher frequency rate amongst specific sub-groups of the overall participant population.

The final step in the process of generating experiential themes was to evaluate the finalised theme in terms of its descriptive validity (how well the theme captured the empirical data)

and its empirical adequacy (was there sufficient data to support the claim being made). This was achieved through re-reading the transcript data, to see how well the finalised theme reflected the originally coded data and whether any examples of the theme had been overlooked. This rereading of the transcript data also provided an additional opportunity to identify examples of participants whose experiences ran counter to the observed theme, and which might, therefore, merit further analysis. In some instances, it was also possible to triangulate the themes against empirical data from other sources, e.g. data contained in the participant reflective exercises.

Development and evaluation of inferential themes

Inferential themes relate to those aspects of the social world that cannot be directly observed empirically, but which can reasonably be inferred through empirical investigation or personal experience. Inferential themes, therefore, reflect what is happening in the ontological domain that critical realists call the *actual*. In contrast to experiential themes, which are descriptive and grounded in the empirical data, inferential themes seek to extend the scope of experiential themes using a combination of inductive and abductive thinking (Wiltshire and Ronkainen, 2021). The aim of inductive reasoning is to move beyond themes that seek to describe the perceptions of participants in the research, to more tentative statements that capture what is being perceived by similar groups of individuals, beyond the participant population. Hence there is a shift from considering the experiences of specific individuals, to more general events (Fryer, 2022). This is achieved using probabilistic language, the precise nature of which should reflect the balance between cases that confirm or disconfirm a given statement. Abductive reasoning is used to move beyond the lay language of participants to the use of more conceptual language that redescribes specific experiences in a more abstracted way.

In my research the development of inferential themes meant moving beyond the experiences of legal professionals within the two firms from which my empirical data was generated, to thinking about what this meant about the experiences of legal professionals more widely. The more abstracted statements produced by this step in the analysis were useful as they offered a more general insight into how AI-enabled legal services might be

understood by legal professionals and, therefore, how AI-enabled legal services can be conceptualised as a phenomenon.

This step of the analysis meant considering each experiential theme in turn, and reflecting on the degree to which it reflected the available empirical data. First, through reference to the data generated by my own research, inferences that could be linked to participant experiences that were more frequent across the population, were granted a greater degree of confidence than inferences where the empirical data was less homogenous. Second, inference strength was also considered through reference to the findings of other studies looking at similar phenomena. Inferences that were supported with evidence from other contexts were regarded as stronger than those for which there was no further data available, or the available data were disconfirming of the claim being made. Where the inference being made could be linked to an existing concept or theory, the extant language used in the literature was drawn upon, in order to reduce ambiguity. In instances where more than one experiential theme was linked to a shared concept, the themes were consolidated into a single inferential theme (Fryer, 2022).

Each inferential theme that was developed from my analysis utilised the suggested stem, “It is plausible to claim that legal professionals [could/often/are likely to] ...” (Wiltshire and Ronkainen, 2021). For example, the experiential theme,

“Most legal professionals do not feel more negatively about their role from using AI-enabled legal services.”

Was reworded through reference to the extant literature relating to how changes in information technology can cause feelings of identity threat amongst professionals (e.g. Jussupow *et al.*, 2018), and became the inferential theme,

“It is plausible to claim that legal professionals are unlikely to experience identity threat from using AI-enabled legal services.”

Evaluation of the inferential themes was based on a combination of their empirical adequacy (was there sufficient data to support the claim being made) and their interpretive

validity (the extent to which my interpretations reflected the perceptions and experiences of the participants).

Development and evaluation of dispositional themes

Dispositional themes reflect the deepest of Bhasker's ontological domains, what he calls the *real*. Dispositional themes can be regarded as attempts to generate a plausible hypothesis, in terms of causal factors and mechanisms, which can explain the empirical outcomes and inferred events identified earlier in the analysis (Wiltshire and Ronkainen, 2021). The themes are generated using retroduction, a specific form of abductive reasoning (Bhaskar, 1998). This required me to reason backwards, asking myself what processes and mechanisms could have produced the outcome of interest. In generating my hypotheses, I drew upon my knowledge of the extant literature, insights within the research data, and my own experiences of using AI and working in a law firm environment. For example, having inferred that identity threat was unlikely to be experienced by legal professionals, as a result of using AI-enabled legal services, retroduction required me to think about what theories or scenarios would best explain this outcome. The resulting dispositional theme in this instance was,

"The inferred phenomenon [an absence of identity threat] is dependent upon the existence of a user-tool perspective of technology amongst legal professionals."

This was based on my evaluation that if legal professionals considered AI to be a tool that they possessed agency over, this would explain why they were unlikely to consider AI as a threat to their professional expertise, status and autonomy.

When more than one plausible theoretical explanation was identified, I evaluated each hypothesis with reference to both the theoretical validity of the explanation (the extent to which the proposed dispositional theme explained the related experiential and inferential themes) and its explanatory power (the extent to which the theme provided a parsimonious explanation, when compared to competing theories). In the above example, this process of evaluation gave me additional confidence that the existence of the technology as a tool

perspective amongst legal professionals was a plausible theoretical mechanism, as it provided a partial explanation for two of the inferential themes identified in the research.

Having completed this final stage of analysis I then presented a summary of the three different levels of theme in diagrammatic form, what Wiltshire and Ronkainen (2021) call an explanatory statement. The aim of this was to organise the connected themes in a way that reflected the proposed causal path that had emerged from the analysis, and which highlighted the links between each theme. This means the order the themes are presented in the explanatory statement reverses the order in which they were derived, with the more abstract themes, which can account for the observed empirical data, presented first. An illustrative example of an explanatory statement derived from the dispositional theme relating a user-tool perspective amongst legal professionals, is provided below (Figure 9).

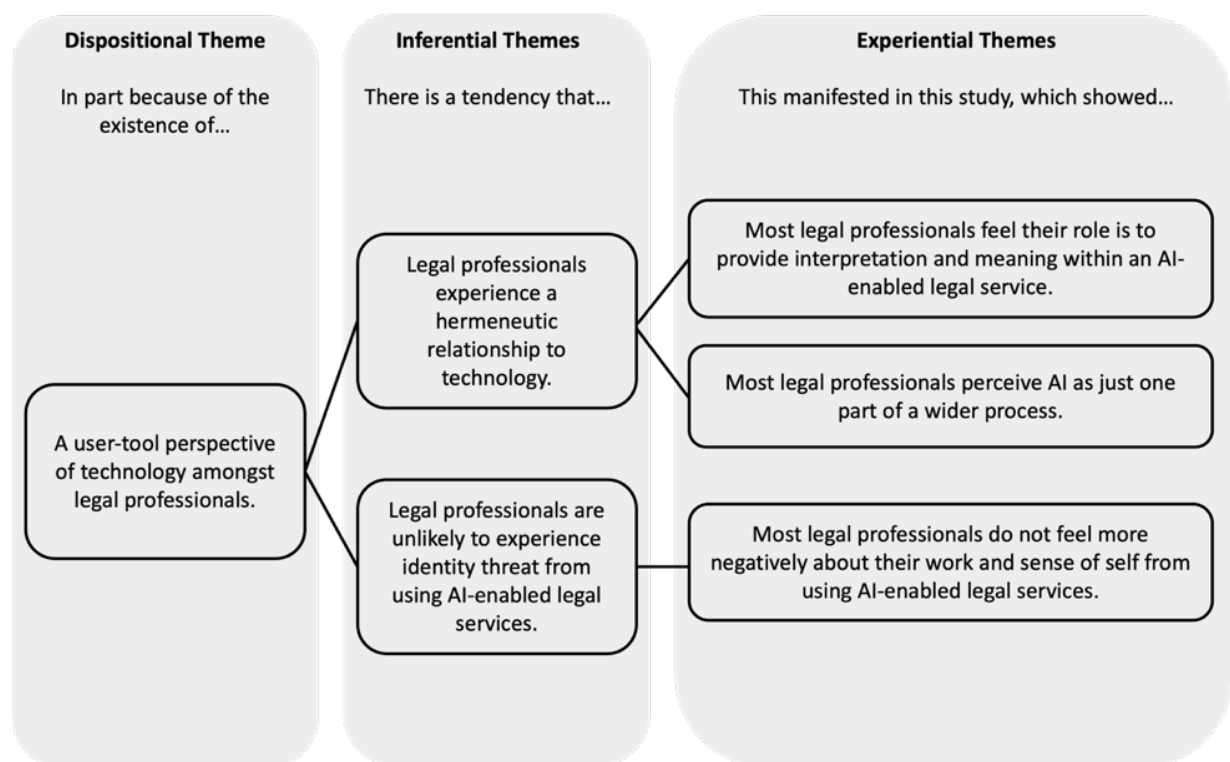


Figure 9: An example of forming explanatory statements from the thematic analysis

Adapted from Wiltshire and Ronkainen (2021)

5.5 Process for identifying causality in case studies

Related research Question: What explained the use of AI-enabled legal services amongst these legal professionals?

Theory building process tracing was used in this analysis to develop the findings of the thematic analysis (discussed in Chapter 6), which highlighted differences amongst *consumers, producers* and *liminals*, relating to the use of AI-enabled legal services.

While the thematic analysis was useful in highlighting cross-case differences, the analysis did not allow for the detailed within-case analysis that is necessary to investigate *how* the use of AI took place. How did the different causal factors interact with one another - which were necessary or sufficient for AI use to take place? What were the causal mechanisms that explained how the identified causes were linked to the use of AI-enabled legal services? This second stage of the analysis sought to trace the causal pathway leading to AI use, to better understand the processes involved. The analysis employed a three-step model of theory building process tracing (Beach and Pedersen, 2016), which is summarised below (Table 17), before being discussed in greater detail.

Table 17: Theory building process tracing

Theory-building process tracing (Beach and Pedersen, 2016)
Step 1 – Selection of typical case A typical case was selected from a bounded set of cases that each contained a similar set of plausible causal factors linked to the outcome of interest.
Step 2 – Development of case narrative A detailed empirical narrative of each typical case was developed; allowing the identification of fingerprints that indicated a specific causal mechanism.
Step 3 – Evaluation of empirical evidence Assessment of the extent to which the identified fingerprints justify the inferred causal mechanism; this is achieved through reference to wider extant theory that could confirm/disconfirm the inference.

Source: Beach and Pedersen (2016)

Selection of typical cases

Before undertaking the more detailed within-case analysis, comparative case methods were used to construct a causally and contextually homogenous population of cases (as a subset of the overall case population), all of which demonstrated the previously identified outcome of interest i.e. the use of AI-enabled legal services by legal professionals. Creating a bounded population of similar cases was important for two-reasons: first, it facilitates the selection of a typical case for detailed within case analysis; second it creates a group of causally similar cases that the findings of the within-case analysis can be generalised to. Generalising findings beyond this bounded population is problematic, as the expectation is that causal factors will have differential outcomes when combined in different ways or observed across contexts (Beach and Pedersen, 2016).

The populations were created through reference to the data generated in the earlier thematic analysis and supporting data. Using a positive-on-outcome comparison meant the case population was initially restricted to only cases where there was evidence that AI-enabled legal services had been accepted – because of the homogeneity of the case population in respect of this outcome this meant all cases were retained at this point.

I then took the decision to remove cases 5, 7 and 13 from the causal analysis as the data relating to these cases was incomplete because these participants did not complete both of the written exercises. While this had not prevented me analysing the different sources of data separately, it would have proven problematic to include them in this stage of the analysis which relied on identifying the causal factors related to each case by considering data from all sources.

Then with reference to my knowledge of the individual cases, developed during the earlier stages of analysis I generated a list of plausible causal factors that could be linked to the use of AI-enabled legal services. The list contained causal factors with powers seen to emerge from the affordances of AI-enabled legal services; the agentic actions of individual legal professionals; and normative powers emerging from larger structural entities including clients, teams and law firms as organisations.

I then asked myself the question, 'Which causal condition(s) is(are) shared by all the cases?' The existence of a single causal condition across all the cases would have suggested that this could be a necessary causal condition for the use of AI-enabled legal services.

A single sufficient cause was not identified, although causal powers linked to AI-enabled legal services; and an individual's team did show coverage across the vast majority of cases. This wide coverage reflected my decision to combine the causal powers relating to a single entity in order to make the process of analysis practical.

I then sought to identify tentative combinations of causes that together might be sufficient to cause AI use. Once this list was established each member of the case population was assessed to determine how many of the causal factors it shared, with the aim of identifying a subset of cases that were maximised for similarity.

A typical case was then identified for each bounded population that was identified. This was based upon a combination of the confidence with which they could be considered a member of the group, which was reflected in the number of causal factors the case shared with group; and the availability of detailed empirical data, which was necessary to perform the process tracing analysis, and which would enhance the reliability of the findings.

Development of case narrative

A case narrative was then developed for the two typical cases selected for process tracing by organising the available empirical evidence for each case. The aim of the narrative was to create a record of who did what and when, in order to facilitate a more systematic case analysis (Beach and Pedersen, 2016). The finalised content of the narrative, therefore, provided me with an initial starting point for my theorizing about the type of mechanism (i.e. the type of activity taking place and the entities involved) that might explain the use of AI-enabled legal services. Given that interaction between legal professionals and AI-enabled legal services is inherent to the process of acceptance, a socio-technical perspective was reflected in the case narratives of both typical cases, although the emphasis on different causal factors was evident.

Through reference to the empirical fingerprints of each case, I then sought to infer whether the causal mechanism was primarily related to the wider social structure in which AI-enabled legal services was embedded; or had emerged from the AI-enabled legal services process itself, potentially reflecting the design and functionality of the service. These fingerprints took different forms, including trace evidence, the existence of which would support a particular explanation, or more broadly, patterns or sequences in the data that would be indicative of a specific theoretical explanation.

This required me to use a retroductive thinking process, similar to that employed when generating the dispositional themes of the thematic analysis. Any potential mechanisms I proposed needed to be able to account for the empirical fingerprints while also linking one or more of the causal factors to the use of AI-enabled legal services. This process led to the generation of two distinct causal pathways that led to AI-enabled legal services use – one for *producers* and one for *consumers*. The theory-building process is illustrated in Figure 10.

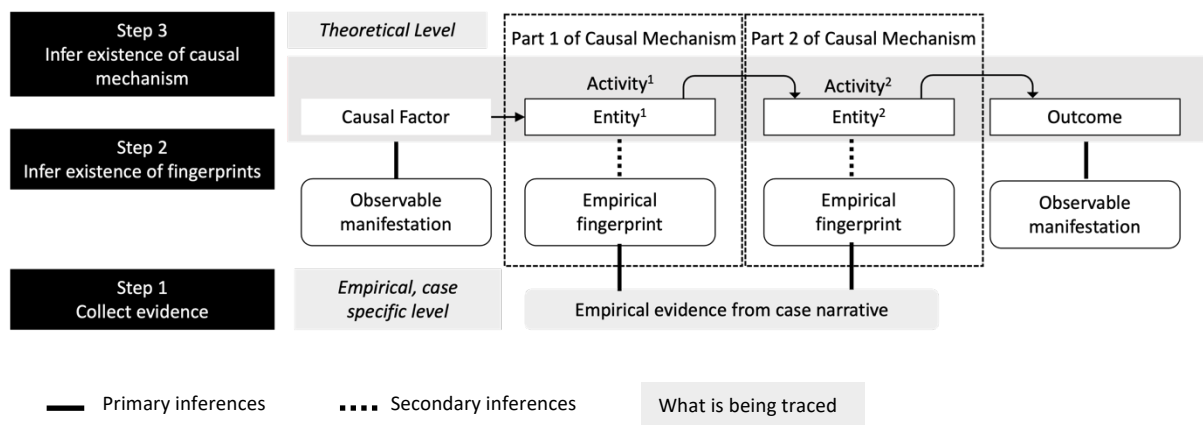


Figure 10: Theory building process tracing

Source: Adapted from Beach and Pedersen (2016)

Evaluation of empirical evidence

Having generated a causal mechanism for each typical case, based on the identified empirical fingerprints, I then sought to evaluate my proposed mechanisms. This involved consulting the extant literature, in order to see whether there was any wider research that would give support to the proposed mechanism; or alternatively whether other competing

mechanisms might offer a superior explanation (i.e. better able to explain the empirical fingerprints) to the one I had generated.

While my research methods were intended to minimise the limitations associated with interviews, my process of evaluation also involved reflecting on the degree to which I should trust the empirical data I generated with my research participants. This involved evaluating the proximity of the participant to the events they described to me, were they directly involved or was their account based on information provided to them by other sources? The limits of human cognition also needed to be taken into account, what cognitive biases might have shaped the participant's account of events, was their memory of past events accurate? Finally, it was important to consider the participant motives for taking part in the research. Thinking about each interview in reference to Alvesson's (2011) eight metaphors of interviews, was helpful in structuring my thoughts. For example, was the participant trying to present themselves or their organisation in a favourable way? Was their account constrained by factors outside their control or reflective of an organisational 'script'?

Having completed a theoretical evaluation of my theoretical mechanisms I then sought to assess its explanatory power by determining the extent to which the mechanism generalised to other similar cases within the original bounded population. This meant looking for the existence of empirical fingerprints similar to those identified in the typical case, from which the mechanism had been originally inferred. Evidence of empirical fingerprints that were theoretically consistent with the proposed causal mechanisms, were regarded as offering cautious support for the presence of the proposed mechanism across other cases.

6. DATA ANALYSIS – THEMATIC ANALYSIS

This chapter presents the findings of the thematic analysis I undertook using data generated from 21 semi-structured electronic interviews undertaken with legal professionals from law firms National and Global. In addition to the electronic interview data, the analysis also makes use of supplementary data from the written exercises that participants were asked to complete prior to their interview; two background interviews with organisational sponsors; and observations I made during demonstrations of the AI-enabled legal services at National and Global.

The chapter begins with a summary of the descriptive statistics of the supplementary data, summarising the contextual data relating to each case and the data relating to the two reflective exercises (section 6.1). I then explain the development of each of the dispositional themes using the five-stage process of analysis outlined in Chapter 5 (section 6.2). Each theme is illustrated with quotes taken from the semi-structured electronic interviews. Finally, drawing on all of the above, I reorganise the data and present it in the format of a technological frame, with the dimensional structure of the frame and the content within each identified dimension, reflecting the knowledge, assumptions and expectations of AI-enabled legal services, held by the participants. While all participants developed individual frames which facilitated their own use of AI-enabled legal services, the data from the thematic analysis made it possible to identify evidence of both frame homogeneity and heterogeneity amongst the participants. Contrasting technological frames are presented for groups of participants identified as producers and consumers of AI-enabled legal services. The presentation of the data in the format of a technological frame is designed to help demonstrate how the findings of the research contribute to answering the research question, *How are AI-enabled legal services understood by the legal professionals that use them?*

6.1 Supplementary Data

During the initial review of the interview data and prior to conducting the thematic analysis I noted contextual data relating to the demographics of each member of the case population and their responses. I also sought to categorise the data within the two reflective exercises, to aid comparison between cases and identify patterns in the responses of the population as a whole. While this data did not form part of the empirical data that was used to generate the initial experiential themes, I did refer to it when developing my interpretations of the data and when evaluating the themes. The contextual data is discussed below, with a summary of the data for each case provided in Table 18.

Contextual data

The law firm each participant worked for at the time of the interview was recorded as either Global (GL) or National (NT); of the 21 participants, 13 worked for Global and 8 for National. The gender of each participant was recorded as Female (F) or Male (M); of the 21 participants, 15 were Female and 6 Male. While gender differences in relation to the use of AI-enabled legal services were not a focus of the study, the higher proportion of female participants reflects the fact that for several years the majority of new entrants to the UK legal professional have been women (Aulakh *et al.*, 2017).

The participant's job role was coded to one of five categories. Trainee Fee Earner (TFE) described individuals working in a fee-earning role but who were not yet qualified solicitors. Qualified Fee Earners (QFE) was used to describe individuals working in fee-earning roles with less than five years of experience following qualification. Experienced Qualified Fee Earners (EQFE) was used to describe individuals working in fee-earning roles with over five years of experience following qualification. Leaders (L) referred to individuals with a leadership role in their team or the wider firm. Individuals with technology-focused roles and without legal fee-earning experience were described as Technical Specialists (TS). Of the 21 participants, there were 3 TFE, 4 QFE, 9 EQFE, 3 L and 2 TS. While analysis of the results within each group was not a focus of the research, one of my methodological goals was to generate data with different stakeholders, to ensure the views of all groups were reflected in the research. The final sample broadly reflects this ambition, with the number of participants in each group reflecting the size of their group as a proportion of the firm.

The current specialism of each participant was also recorded to one of four categories. Fee-earners whose work was primarily of a contentious nature, for example general litigators and those specialising in professional negligence claims, were described as Contentious Lawyers (CL). Fee-earners whose legal specialism focused on transactional work e.g. property solicitors and banking solicitors were described as Non-Contentious Lawyers (NCL). Individuals working in teams focused on the development and deployment of AI-enabled legal services, were described as having a Multi-Disciplinary (MD) specialism. Trainee Fee Earners were not allocated a specialism (N/A) as their role required them to undertake both contentious and non-contentious work. Of the 21 participants, there were 4 CL, 7 NCL, 7 MD and 3 N/A. When compared to the overall population of these groups within each firm, it is important to note that individuals working within multi-disciplinary teams were over-represented in the sample population. This should be taken into account when evaluating whose views the analysis can be seen to reflect.

The year in which each participant first started using AI-enabled legal services was also recorded. Responses ranged from 2010 to 2021, with most participants (16) reporting they had started using AI-enabled legal services between 2017-2021. This indicates that most participants have only been using AI-enabled legal services for a relatively limited period, which reflects the recent growth of the phenomenon. However, it is also important to remember that the majority of legal professionals in the UK are not currently using AI-enabled legal services, meaning the participants in this study are relatively speaking experienced users.

The way in which each participant interacted with AI-enabled legal services through their work was also recorded in one of three categories. The first two categories reflected Armour *et al.*'s (2022) distinction between *Consumers* and *Producers* of legal services. *Consumers* refer to individuals whose work is augmented by AI technology but who do not have a direct role in the development of AI-enabled legal services. This is an accurate description of most fee-earners in a law firm. In contrast, *Producers* are persons with legal human capital working in a multi-disciplinary team (MDT) that is responsible for the technology that augments AI-enabled legal services. Using Armour *et al.*'s definitions I was able to clearly categorise 10 Consumers and 5 Producers within the overall population.

However, there were a further 6 individuals whose interactions with AI-enabled legal services made it difficult to categorise them as either producers or consumers. These were individuals who had experience in a producer role, typically as a result of being temporarily seconded to a multi-disciplinary team, before returning to their usual fee-earning work. At Global the secondments were full-time typically for a period of between 6-12 months.

“So I've been at Global since I qualified as a solicitor, so way back when in 2013, erm and I got recruited into an internal secondment into the [Global's multi-disciplinary team] in February 2022...I handed over most of my transactional fee-earning stuff back in February, and I've been doing this exclusively...it wasn't intended to be too structured, we'd look at what we were doing; we would evolve our approach based on how different things were happening and changes we were seeing, and it would just be a very loose fairly unstructured role, which I think is just sort of the nature of what you're doing.” [Participant 19]

At National secondments were aligned to an AI development project, with individuals working as a producer one day per week for the duration of the project, alongside their normal fee-earning duties.

“I've been a National for 10 years in October, believe it or not. And yeah, I always worked in this [medical claims] team more or less and always sort of done this work and, to be honest, worked with the same clients for most of those 10 years...Well, at the moment we use and we're trialling software to do with reserves. And so that's what we're doing at the moment that uses AI, uses it, uses AI I think to a certain extent, it's sort of a predictor. ... That's the project that I'm involved in at the moment specifically. And then there've been other projects too...a prototype which is been tweaked and sort of that level, you know, it's not day to day.” [Participant 5].

I decided to categorise these individuals as *Liminal* (LIM), reflecting the boundary-spanning nature of the work undertaken by these individuals, which meant they performed tasks that were characteristic of both producers and consumers of legal services, and yet could not be accurately described as members of either of these groups. The fact that these liminal roles were being performed on a temporary and/or part-time basis, also seemed to reflect the ambiguous status associated with liminal experiences.

Table 18: Contextual data relating to each participant

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
Law Firm	NT	NT	NT	NT	NT	NT	NT	NT	GL	GL	GL	GL	GL	GL	GL	GL	GL	GL	GL	GL	GL
Gender	F	F	F	M	F	F	M	F	F	M	F	F	M	F	F	M	F	M	F	F	F
Job role	L	TFE	TFE	EQFE	EQFE	TS	L	TFE	EQFE	EQFE	QFE	QFE	TS	EQFE	L	QFE	EQFE	EQFE	EQFE	EQFE	QFE
Specialism	MD	N/A	N/A	NCL	CL	MD	NCL	N/A	CL	CL	CL	MD	MD	NCL	NCL	NCL	MD	NCL	MD	MD	NCL
Experience of AI-enabled legal services	2017	2021	2018	2020	2021	2021	2017	2019	2019	2017	2019	2020	2010	2017	2015	2017	2017	2015	2012	2015	2021
Interaction with AI-enabled legal services	PR	LIM	CS	CS	LIM	PR	CS	LIM	CS	CS	CS	PR	PR	CS	CS	CS	PR	CS	LIM	LIM	LIM

Table abbreviations

Law Firm: GL, Global LLP; NT, National LLP

Gender: F, female; M, male

Job role: EQFE, experienced qualified fee earner; L, leader; QFE, qualified fee earner; TFE, trainee fee earner; TS, technical specialist

Specialism: CL, contentious law; MD, multi-disciplinary team; N/A, not applicable; NCL, non-contentious law

Interaction with AI-enabled legal services: CS, consumer; LIM, liminal role; PR, producer

Reflective exercise data

Reflective exercise one: Narrative timeline

A narrative timeline was completed by 19 of the 21 participants in the study prior to their interview (participants 5 and 14 did not complete the exercise); a summary of the data is shown in Table 19. The narrative timeline exercise captured the experiences that the participant identified as being relevant to the formation of their attitude towards AI-enabled legal services. In all instances participants reported a mixture of positive and negative experiences of varying intensity. The intensity of the experience is indicated as being High (H), Medium (M) and Low (L), or NR where no rating of intensity was provided. While each participant recorded a unique combination of experiences in their timeline, certain experiences were reported more frequently than others. Experiences involving interactions with other people featured most frequently, although first-hand experiences that involved only the participant also featured.

The most frequently cited positive experience (18 of the 19 participants) was *being encouraged to use AI-enabled legal services by their wider team*, the intensity of these experiences also tended to be highly positive. In contrast, *positive client feedback about AI-enabled legal services* only featured in the accounts of seven participants. Less experienced participants who had more recently entered the legal profession, were more likely to highlight the positive impact of experiences relating to AI-enabled legal services *during their academic studies*.

Negative experiences were reported less frequently than positive experiences, but three negative experiences were reported by a significant minority of participants. *Resistance/disinterest towards AI-enabled legal services amongst colleagues* was reported by nine participants. Negative user experiences also featured prominently with nine participants indicating *AI-enabled legal services had not performed as expected*, and eight participants indicated *AI-enabled legal services had not been able to meet their specific needs*.

The timeline data also revealed differences in the responses of *consumers, producers* and *liminals*. On average, producers and liminals reported a wider range of positive experiences

affecting their attitude towards AI-enabled legal services, than consumers. *Positive client feedback* was cited significantly less frequently by consumers (2 of the 9) compared to half of those in producer/liminal roles (5 of the 10). *Working with external providers of AI* was not mentioned by any consumers but had a highly positive impact on half of those in producer and liminal roles (5 of the 10). Most liminals also highlighted being *seconded to a multidisciplinary team* as being highly positive; an experience that was not reported by consumers or producers (the latter being permanent members of multidisciplinary teams). Producers were the only group to report *delivering AI projects for the firm* as highly positive in significant numbers (3 of 4).

The pattern for negative experience between the groups also revealed important differences. *Resistance/disinterest of colleagues* was far less frequently reported by consumers (3 of 9) compared to a majority of producers and liminals (6 of 10). Interestingly, *AI not working as anticipated* was mentioned by the majority of liminals (5 of 6) but none of the producers (0 of 4), which may reflect the liminals' experiences as consumers when undertaking fee-earning work.

Table 19: Narrative timeline data relating to each participant

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
Positive Experience																					
AI use encouraged by wider team		M	H	H	H	H	N/A	H	H	H	NR	M	N/A	H	H	M	H	H	M	NR	H
AI received positive client feedback	H				H	H	N/A				NR		N/A		H				M		M
Use of AI discussed while studying		M	H	H		H	N/A	M					N/A					M			
Secondment to MDT developing AI		H					N/A	H				H	N/A						H		M
Working with external providers of AI	H					H	N/A	H					N/A				H		H		
External recognition for use of AI	H						N/A					M	N/A	M					M		
Delivering AI projects for firm	H						N/A					H	N/A				H				M
Attending LegalTech events						L	N/A						N/A			H				NR	H
Negative Experience																					
Resistance/disinterest amongst colleagues	M			M			N/A	L					N/A			H	H	M	M	NR	H
AI did not perform as expected		M	M		M		N/A	M		M			N/A		M	M			L		M
AI could not meet specific user needs		M		M		M	N/A			H	NR		N/A			M			L		
AI use not encouraged by wider team							N/A				NR	L	N/A							NR	
Resistance from clients	M			H			N/A						N/A	M							

Table abbreviations

Intensity of experience: H, high; L, low; M, medium; NR, intensity not recorded

N/A, participant did not provide data

Reflective exercise two: How AI affects your work

The tetrad exercise indicating how AI affected their work, was completed by 18 of the 21 participants in the study prior to their interview (participants 5, 10 and 14 did not complete the exercise). The tetrad used the four laws of media (McLuhan and McLuhan, 1988) as a framework to capture how the participant understood AI-enabled legal services to have affected their work. A summary of the data is shown in Table 20; statements that were reflected in the candidate's completed exercise are recorded as 'True'. If the statement was not mentioned the field is left blank, but this does not indicate disagreement with the statement.

Despite being provided with a worked example of the exercise in their materials, participants reported during the interview that they had found this exercise more difficult to complete than the narrative timeline. This led to inconsistencies in the level of detail provided, with some participants providing a much higher level of detail than others. This was most apparent in relation to the third law of media ('Retrieves') which asked candidates to indicate, *What previously obsolete tasks or capacities has your workplace AI made possible again?* In interpreting the results I was therefore more cautious about making claims about what the data could be seen to indicate.

Most (17 of 18) participants indicated that AI-enabled legal services had *enhanced* (Law of media 1) their work in some way. With enhancements in terms of the *speed* (13 participants), *accuracy* (12) and *quality* (10) of work highlighted by a majority of the participants. The ability of AI-enabled legal services to enhance the completion of high-volume work was highlighted by 8 participants.

Most (17 of 18) participants also recognised that AI-enabled legal services had *diminished* (Law of media 2) aspects of their work. This included the *development of their core skills* (11 participants) and some of the more *administrative aspects* of their work (8). It is important to note that diminished is understood to be a value neutral term, it is meant to reflect tasks or actions that have been reduced or made obsolete. Whether participants understood these changes to be positive or negative, was explored later in the interview. This allowed me to understand that while lawyers saw the diminishing of their core skills as a negative

development, reductions in the volume of administrative tasks they were expected to perform was welcomed.

AI-enabled legal services were also recognised to have *retrieved* (Law of media 3) the ability to review all the data relating to a specific legal matter by a minority of the participants (5 of 19); something that had not been possible when the reviewing needed to be conducted by humans.

When asked to think about how AI-enabled legal services might impact their work if used to their maximum potential (Law 4 – ‘*reverses into*’), most participants (16 of 18) provided a response. They identified a number of potential outcomes, most of which were linked to economic performance and efficiency. The *deskilling of legal professionals* was the most frequently cited outcome (7 participants) with a *reduced need for legal professionals* (6) and *support staff* (4) also highlighted.

The tetrad data also highlighted certain differences in the ways in which *producers*, *consumers* and those in *liminal* roles perceived AI to have affected their work; this is not surprising given the work they undertake varies considerably according to the precise requirements of their role. While a majority of consumers (5 of 9) highlight AI-enabled legal services as enhancing their *ability to work at scale*, this was mentioned by only a minority of producers/liminals. This may reflect the fact that their roles tend to be project based and less likely to involve the type of high-volume tasks fee earners undertake. Perhaps unsurprisingly, it was those with fee-earning responsibilities (a majority of both consumers and liminals) that indicated AI was diminishing the *development of core skills*. Interestingly while half of those in producer/liminal roles saw AI reducing the need for lawyers, only 1 of 9 consumers highlighted this as a concern. This would seem to counter discourses emanating from the future of work debate, which suggest a significant number of professional roles are at risk of technological substitution.

Table 20: Tetrad exercise data relating to each participant

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	
Enhances																						
Speed	TRUE	TRUE	TRUE	TRUE	N/A	TRUE	N/A	TRUE				TRUE	N/A	TRUE			TRUE	TRUE	TRUE	TRUE	TRUE	TRUE
Accuracy	TRUE		TRUE	TRUE	N/A	TRUE	N/A			TRUE	TRUE	TRUE	N/A	TRUE		TRUE			TRUE	TRUE	TRUE	TRUE
Quality					N/A	TRUE	N/A	TRUE			TRUE	TRUE	N/A		TRUE	TRUE	TRUE	TRUE	TRUE	TRUE		TRUE
Working at scale		TRUE		TRUE	N/A		N/A						N/A	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE		
Use of historic data					N/A		N/A						N/A		TRUE		TRUE	TRUE				
Obsolesces or Diminishes																						
Development of core skills	TRUE	TRUE			N/A		N/A	TRUE			TRUE	TRUE	N/A	TRUE	TRUE	TRUE		TRUE	TRUE		TRUE	TRUE
Administrative tasks				TRUE	N/A	TRUE	N/A			TRUE			N/A		TRUE		TRUE	TRUE	TRUE	TRUE		
Creation of bespoke solutions					N/A		N/A		TRUE			TRUE	N/A						TRUE			
Interaction amongst colleagues					N/A		N/A		TRUE				N/A	TRUE		TRUE						
Retrieves																						
Ability to review all data					N/A	TRUE	N/A				TRUE		N/A			TRUE				TRUE	TRUE	TRUE
Reverses Into																						
Deskilling of lawyers				TRUE	N/A		N/A	TRUE	TRUE				N/A	TRUE		TRUE			TRUE		TRUE	TRUE
Reduced need for lawyers	TRUE				N/A	TRUE	N/A						N/A				TRUE	TRUE	TRUE		TRUE	TRUE
Professionals focus only on the complex		TRUE			N/A	TRUE	N/A				TRUE		N/A					TRUE	TRUE			
Overarching focus on efficiency					N/A		N/A						N/A		TRUE			TRUE	TRUE	TRUE		
Enhanced profitability		TRUE			N/A		N/A						N/A		TRUE			TRUE				TRUE
Reduced need for support staff					N/A		N/A			TRUE			N/A			TRUE		TRUE	TRUE			
Automation of non-legal tasks		TRUE			N/A	TRUE	N/A						N/A			TRUE						
Over-reliance on technology				TRUE	N/A		N/A						N/A						TRUE	TRUE		

Background interviews and observations

Data generated with the organisational sponsors at National and Global gave me the opportunity to assess the extent to which AI-enabled legal services at each firm conformed to the tool perspective of technology. My evaluation of each AI-enabled service was based on a combination of: technical information about the AI-enabled services at the firm and how they had been implemented, provided by the organisational sponsor; my own observations of the AI-enabled services, when they were demonstrated to me by the organisational sponsor; and publicly available information from the AI software developers (when the firm had chosen an externally developed AI technology solution, to underpin their services).

My evaluation process involved using the above information to determine which (if any) of the five assumptions identified by Schuetz and Venkatesh (2020) were being challenged by the AI-enabled legal service (Tables 21 and 22). My analysis showed that when the AI technology underpinning the service had been developed within the firm, it conformed to the five assumptions of the tool perspective. This reflected deliberate choices made by each firm to develop AI-enabled services that could be explained and justified to lawyers and clients of the firm (if required); and to allow individuals to choose whether to use AI, rather than make its use mandatory. Relaxing any of the five assumptions would have undermined the achievement of these goals. However, both organisational sponsors accepted that many of the individuals that used AI-enabled legal services would not have a detailed technical understanding of how the AI technology underpinning the service worked. These individuals might, therefore, perceive AI as a 'black box', even though other members of the firm possessed this technical knowledge.

When AI-enabled services were underpinned by AI technologies developed by outside vendors, it was more difficult to judge whether the five assumptions of the tool perspective were being challenged. This was because the organisational sponsors at National and Global relied on technical information provided by the external AI developers to understand how the AI worked. My evaluation of the systems has, therefore, assumed this information to be accurate and complete. This led me to conclude that two of the three external

systems conformed to the tool perspective, but one system (System Z) appeared to exhibit characteristics that were inconsistent with the assumptions of environmental ignorance and functional consistency.

AI-enabled legal services at National

National highlighted three different AI-enabled legal services at the firm. Services A and B were both designed to make claims forecasting more consistent and were both developed internally at the firm. The technology that underpins these services is best described as an expert system, which employ rules-based AI, to generate an output. The rules were developed through collaboration between the firm's IT developers and legal subject matter experts with access to the firm's historic claims data. This means these experts both determined and understood each of the rules used by the system. The rules do not change over time unless the firm decides to update them, which requires further human intervention. Once implemented, the system required individual lawyers to input specific data relating to their client's claim into a smart questionnaire whose questions are conditional on the previous answers given. The system provides a recommendation based on the information provided, however, lawyers are not required to follow the recommendation, it is just one element to consider as part of their wider decision-making about the claim. The output generated by the system was expected to be broadly consistent with the response a subject matter expert would have given if presented with the same information.

Service C was designed to review contracts and identify potential anomalies, the aim being to improve the speed and accuracy of due diligence activities at the firm. The system was underpinned by AI software from an external vendor, which was then tailored to the needs of the firm. The software used machine learning to develop algorithmic decision-making models capable of identifying relevant information. The software required 'training' to achieve this, which meant subject matter experts from the firm were able to oversee the development of the templates the system used to review documents; the templates were then fixed and could only be changed by authorised experts at the firm. Once implemented, individual lawyers uploaded relevant documents to the system to be reviewed (often several hundred at once). The system then flagged documents that deviated from the

predetermined template for the lawyer to review, before deciding whether any further action needed to be taken.

Table 21: Evaluation of AI-enabled legal services at National

	Service A	Service B	Service C
Function	Claims forecasting	Claims forecasting	Document review
AI Development	Internal	Internal	External
1. Unilateral relationship	Yes	Yes	Yes
2. Ignorance of environment	Yes	Yes	Yes
3. Functional consistency	Yes	Yes	Yes
4. Functional transparency	Yes	Yes	Yes
5. Awareness of use	Yes	Yes	Yes

AI-enabled legal services at Global

Global highlighted four different AI-enabled legal services at the firm. Services W and Y were employed in different parts of the firm, but both were designed to improve the speed and accuracy of due diligence activities through using AI to ‘tag’ specific data within documents. Although the precise methods used by each system were different, they both followed a pattern similar to that employed by Service C at National.

Service X was developed by Global to analyse large volumes of unstructured data and identify patterns in the data that would not be visible to a human legal professional. This allowed the system to draw inferences from existing data, thereby creating new information, which could be used by the firm’s legal professionals. The technology underpinning service X was a rules-based semantic reasoning system, which employed a smart questionnaire, to analyse data inputted by legal professionals. While this superficially sounds similar to System A and B at National, the complexity of the questionnaires used by service X (and the data being inputted) was much higher, with bespoke questionnaires needing to be developed to provide insights to specific client problems. The system was, therefore, transparent in the sense that the firm dictated what rules the system applied, however, the outputs that were produced were ones that could not be replicated by a human legal professional.

Service Z was designed to enhance the ability of legal professionals at Global to forecast case outcomes. The nature of the AI technology that underpinned the system was not known to my organisational sponsor or available from public sources. Hence while the AI may be functionally transparent to the external developers, it was not transparent to Global. This meant Service Z by utilising this AI technology was not transparent to its users. Publicly available information confirmed that the AI technology tracked newly published case data on a daily basis, meaning the data set on which the AI runs was highly dynamic and not visible or under the control of Global. System Z was, therefore, not ignorant of its environment and would not display functional consistency over time.

Table 22: Evaluation of AI-enabled legal services at Global

	Service W	Service X	Service Y	Service Z
Function	Document review	Data Insights	Document review	Case forecasting
AI Development	Internal	Internal	External	External
1. Unilateral relationship	Yes	Yes	Yes	Yes
2. Ignorance of environment	Yes	Yes	Yes	No
3. Functional consistency	Yes	Yes	Yes	No
4. Functional transparency	Yes	Yes	Yes	No
5. Awareness of use	Yes	Yes	Yes	Yes

6.2 Dispositional theme development

Given the embryonic state of research into AI-enabled legal services, the over-arching purpose of the thematic analysis was to get a general understanding of how legal professionals made sense of AI-enabled legal services. This meant starting with a detailed review of the entire data set, rather than selectively focusing on a specific aspect of the participants' experiences. While my analysis was guided by my refined research question and the notes in my 'thoughts and questions' memo, the first step of the thematic analysis used an inductive approach to coding the data, rather than trying to fit data into a pre-existing framework, based on the extant literature or my own views about the topic. The experiential themes that this analysis produced were then used as a starting point, to infer what might be happening more generally, with these ideas being captured as inferential themes. The analysis then concluded with the generation of dispositional themes, where

my interpretation of the data became more abstracted as I drew upon my academic knowledge of the topic and my professional experience within the legal sector.

In this section I discuss each dispositional theme in turn, starting with an explanatory statement that summarises the dispositional theme and shows the causal relationship to its underlying inferential and experiential themes (Figures 11, 12 and 13). The statement presents the themes in the opposite order to which they were generated, as this better captures the direction of causality between the themes, with dispositional themes seen to provide a theoretical explanation for the experiential themes which were grounded in the empirical data.

I then explain the process through which I generated the explanatory statement. Starting with the generation of experiential themes, I provide excerpts from the participant interviews to demonstrate how the experiential themes were grounded in the empirical data. I also indicate where I used the supplementary data to develop and evaluate the experiential themes. I then discuss the inferences I drew from the experiential themes, which were captured as inferential themes; highlighting where the extant literature was used to support my inferences. Finally, I explain which theories I identified as being best able to explain the inferential themes I generated. Where relevant, I indicate the other theoretical explanations I considered, but which were evaluated to be less capable of explaining the empirical data in the study.

A summary of the experiential themes linked to each dispositional theme is presented in Appendix 6. This indicates the specific cases that provided the empirical data from which each experiential theme was developed.

Dispositional theme one: A user-tool perspective amongst legal professionals.

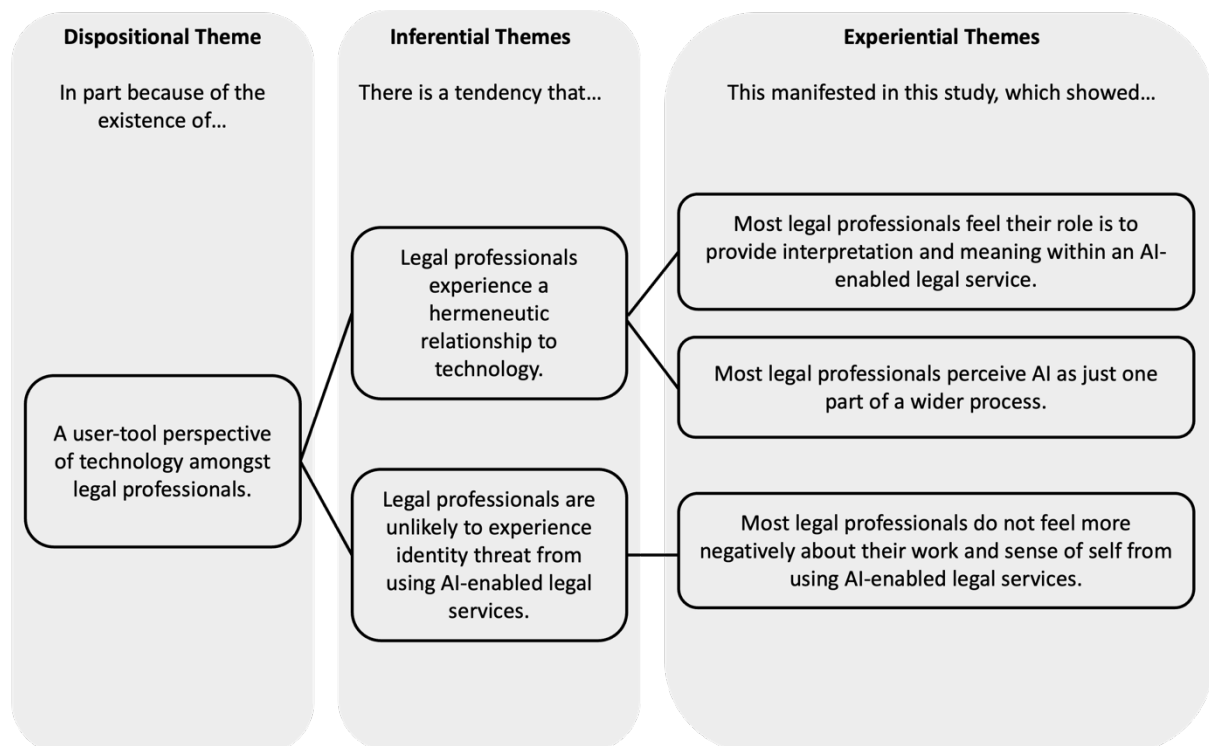


Figure 11: Explanatory statement of dispositional theme one

Experiential themes

1. The data show that most legal professionals in this study feel their role is to provide interpretation and meaning within an AI-enabled legal service.

To understand how legal professionals understood the way in which their relationship to AI-enabled legal services mediated their interactions with the world I asked them to respond to the following question, which used Ihde's (1990) framework for thinking about human-technology relationships.

“What I'd like to do is describe four common ways in which people often talk about relating to technology. What I'm interested to know is if any one of them, or more than one of them, describes how you would relate to the sorts of tools we've been discussing today or whether you discuss your relationship to these tools a bit differently. The first way people sometimes talk about their relationship with the technology is seeing it as an extension of their human body. So, we might think about using a pen to write or a car to drive as being an example of

that. Alternatively, people sometimes talk about a hermeneutic relationship, where you read the technology in order to generate meaning. So, an example of that might be something like a thermometer or a map - the meaning is created by us in response to what the technology presents to us. You've then got the idea of technology as something that's very separate from us. So, like a robot we're interacting with it; potentially we don't necessarily entirely understand it or its behaviour, and we might think it's something that's quite hard to predict. The final idea is that the technology disappears from our conscious view, while remaining in contact with us. An example could be something like a heating system or CCTV camera. Do any of these ideas help explain how you see yourself relating to the tools we've been discussing?"

In response to this question most (18 of the 21) legal professionals indicated that they understood their primary role within the wider AI-enabled legal services as an interpretive one. They described the AI-generated outputs of the process not as the finished legal product, but as information that they transformed through the application of their own legal knowledge and judgement to create a valued service for their client.

"It's probably hermeneutic in that it's a tool that you're using to interpret the data that you are trying to process. I see the extension almost the anthropomorphic extension that you were talking about to be more something you know, like glasses. In that respect, it's kind of more of an object sort of thing, so I think my experience has been hermeneutic in that you're kind of presented with an output, if you will, that you've got to process yourself. And I don't think it's at the point, it's not slick enough for me to be at the point where you're not noticing it." [Participant 4]

"I think the second one [hermeneutic relationship] mostly because we send it away to do something for us, and it will give an output. But then the output doesn't mean anything until we interpret it or, yeah, reject it. I don't use anything that I feel I don't have power over or influence over, and again of all the programs that I use very few of them are just ticking along in the background without me striking a button or selecting an option or hitting a tab." [Participant 9]

“I think it's something that gives you kind of outputs. But again, how you apply these outputs, in some ways even how you apply the inputs as well even, is something that requires a kind of human assessment at the moment. And I think that's to do with the fact that lawyering is not really a, you know, sort of a linear equation, it's there's lots of inputs that are coloured by sort of judgement, particularly things like in a commercial sense. So if you know what's happening in the sector, what competitors are doing, you know what the client has on their horizon. You know, all those kinds of things that will affect on the inputs into the tool and how you assess and or how you assess any output.” [Participant 16]

While most legal professionals characterised their relationship with technology as hermeneutic in nature, other types of relationship were acknowledged. Some (six) legal professionals highlighted relating to some aspects of AI-enabled legal services in an embodied way (first example given in the question). These individuals were all relatively experienced users of AI-enabled legal services, whose role meant they were better understood as *producers* or *liminals* (3 of each), rather than *consumers*, of AI-enabled legal services.

“I think I always think of technology as an extension of yourself and kind of like an ally. It's something to lean on and to use and to utilize, but it's not something that should dictate your actions, necessarily. You can use it to help guide you and to make things possible, and kind of the best way to describe it is when you're working on a Word document, you wouldn't sit there and wait for it to write itself. You have to write it yourself but Word is there to provide the format for you to be able to do it and allows you to be able to write it and type and make it so it is legible and not your horrible handwriting. So, yeah, the embodiment one, I think is the best way to describe it. [Participant 6]

“I'm so familiar with some of the systems I can do it in my sleep. [named AI software] being one, so I'm kind of moving more towards the embodied there, and I completely trust it and I know how it works and I will fully stand behind it if challenged on the output. Others, particularly with some of the ML stuff that we're doing now, which is definitely more of a learning curve for me, a little bit more of a black box.” [Participant 17]

Some legal professionals indicated that their relationship with technology was dynamic or that the relationship varied depending on the specifics of the AI-enabled legal services. They described their relationship changing as their familiarity and experience of the AI-enabled legal services increased. Initially a more separate relationship was experienced (third example given in the question), but over time embodied relationships sometimes developed. There was also speculation that in the future the technology might ultimately disappear from conscious thought i.e. a background relationship (the fourth example given in the question).

“The tools that we've been discussing, I would say, feel like the inter-relationship with a robot kind of one. I feel like it's very much, you know, I need to go, I need to use the thing. Here we go. I'll use the thing. Thank you. Thank you thing. And then you move on from that. I feel it's starting to get a little bit more [hermeneutic]. Yeah, something closer to that, we're somewhere on the cusp of that, where there are things like, you know, the sixth sense coming from, you know, from your laptop...it's not like quite having a chip in your mind yet, but you can see where you can, where you work as one with it, that you're kind of moving towards that sort of final stage, I guess.” [Participant 10]

“I guess longer term, it could probably be any of these, if I'm honest. I think what we're trying to do at the moment is probably something that is embodied, so we're trying to provide something that they actively consciously use for a particular reason. But I can see scenarios where because as something matures and becomes more embedded, it probably satisfies more types of relationships. Some that people won't even be aware that are going on. But then, you know, some of them will continue to be embodied in their use. So, I can see a place where they exist, all these relationships, I think.” [Participant 13]

“I would say for things that have been embedded in our day-to-day work for a long time, like [named AI software (1)], for example. I would say it's one you just you wouldn't even think about it, that's how I would expect to draft a document now. And I think [named AI software (2)] will get to that point. Erm it's new to us in the firm, but I think going forward, people will just always use it to do their Title acquisition. You know, find the documents and review. But I also feel very wedded to two [hermeneutic relationship] in the sense of [named AI software (1)] will take you so far... And I think the bit that lawyers can provide on top of what

technology can do is an interpretation of what that brings out and an understanding of what the client needs from that and bringing the two together, so interpreting what's there and interpreting what the client needs and merging the two to give a good bit of advice. Something like [named AI software (2)] will never be able to do that. [Participant 20]

In evaluating this theme, I felt confident that it provided an accurate reflection of the empirical data. In terms of descriptive validity, 'interpret' was the word most frequently used by participants to describe their role within an AI-enabled legal services, justifying its inclusion in the theme. 'Interpret' was also the word I would have chosen to describe the role of a legal professionals, based on the observations I made when having the AI-enabled legal services demonstrated to me. The empirical adequacy of the theme was supported by the consistency of the participant responses, with 18 of the 21 participants mentioning interpretation as being an aspect of their relationship to AI-enabled legal services, thus justifying the reference in the theme to 'most legal professionals in this study'.

2. The data show that most legal professionals in this study perceive AI as just one part of a wider process.

To develop my understanding of the different elements within an AI-enabled legal services, I asked each participant to explain where the AI technology fitted into the overall process, and what wider practices were linked to the technology. In response to this question, most legal professionals (11 participants) highlighted that AI was one element within a more complex socio-technical process involving other technology and legal professionals. This suggested that they understood AI-enabled legal services as evidence of task augmentation (Rai *et al.*, 2018), rather than task substitution (automation). There was also an emphasis on AI augmenting the work of legal professional, with AI involved in the initial stages of the process, that was being managed and ultimately concluded through the work of legal professionals.

"A client would be acquiring a property that would be subject to a lot of leases. We would want to run those leases through [named AI software] because it would be the quickest way of doing it...Colleagues would be helping me review it, sort of lower-level colleagues, trainees, etc. would be assisting in the review they'd be pulling things out, extrapolating

where things have changed between the leases, why we would need to flag something up, etc. We'd be inputting this information into the reports. And the aim would be that that would then output into a spreadsheet, which we could interpret and, you know, make a bit more client friendly. So that's kind of, in a nutshell, the process that we would adopt and it's [AI] fairly central to the role that we're playing for the client.” [Participant 4]

“There are times where we are doing a review, you know, for example, if we are doing a security review as the kind of bread and butter of a restructuring and insolvency lawyer, sometimes we will be instructed...where you've got maybe 20 documents to review and what we would do is we would use some of that technology, which clusters documents, so groups them together. And all that does, is it shows us the similarities between the documents and it generates comparisons between them all. And so that is still quite a manual job, but that's just where tech kind of assists the project.” [Participant 18]

“I think the technology is a component in what we're doing, so particularly with the title extraction tool for example, it's doing the initial review and extracting the information for us, but it's not adding in any of the legal advice. So, our process would stay the same, we'd still be acting and interacting with clients in the same way that we always had, but we'd have the tool running in the background, probably quite early on in the process. To just weed through the volume of information really quickly, pull out what was key and then use that to then feed into like our standard reports...At the moment, it's just one component that still has to be sort of reviewed by the lawyers quite heavily.” [Participant 19]

In some instances, AI's role in the process was acknowledged to be critical, despite being a relatively small part of the overall process. This stemmed from the fact that the role of AI could not be performed by other actors, meaning without AI either the overall process could not exist, or that it would function in fundamentally different ways. This highlights the importance of not determining the relative importance of different elements of a process according to their magnitude.

“Running the [AI element of the] process itself and doing the process is really small, it's about a five-minute job and you can think nothing else of it, but the answers that we got from that, change the course of the case completely. So, when you are deciding what kind of

strategy to run, when you're reporting back to the client on what your advice is, it is ultimately based on that. So, it is a huge impact but very easily done.” [Participant 2]

In evaluating this theme, I found myself relying more heavily on my own observations of AI-enabled legal services. While most legal professionals did indicate that AI was just part of a wider process, not all participants had enough knowledge of the role of technology to comment on its relative importance or how it interfaced with other aspects of the process, beyond those they were personally involved in. However, the background interviews I held with organisational sponsors, who also demonstrated the use of AI technology to me, provided additional data to support the view that AI-enabled legal services are socio-technical processes in which technology tends to augment the role of legal professionals, primarily through performing a diagnostic role through the processing of data towards the start of the process.

3. The data show that most legal professionals in this study do not feel more negatively about their work or sense of self from using AI-enabled legal services.

To develop my understanding of the ways in which legal professionals understood their work and sense of self to have been changed by AI-enabled legal services, I concluded each interview by asking participants to reflect on whether they felt AI-enabled legal services had impacted their sense of professional identity in any way. In responding to this question all participants (21 participants) indicated the overall impact was either positive or neutral, with no participants reporting their own professional identity as being negatively impacted by AI-enabled legal services.

Legal professionals working in fee-earning roles explained that they did not believe that AI-enabled legal services had led to their role being diminished or replaced. Instead, they highlighted that AI-enabled legal services could actually reduce the administrative burden of their work, and allow them to focus on where they could really make a difference for clients. The introduction of AI-enabled legal services had, therefore, had limited impact on how they felt about themselves in the context of their work.

“You always go into a career as an apprentice with the stereotypes of you'll be doing the filing and the coffee runs. And with AI and technology in the firm that instantly diminishes, that you're not going to be stuck in a corner with a filing cabinet going through pieces of paper shredding and photocopying. Because it doesn't exist for us, you're instantly going to be forced to be doing more the kind of thing we call fee earning tasks.” [Participant 2]

“I think the concern is that, you know, if you if you're taking AI at its basis, it's like, I say, if it's a science fiction thing, it's going to replace you as a lawyer or whatever. I don't think it can do that in its current form. So, I think there's still, you know, there's absolutely a role there it's just how we use how you use it, you know? Does the fact that you don't necessarily have to trawl through reams and reams of paperwork make you less of a lawyer? I wouldn't necessarily say so...Personally, I'm more about, you know, finding the solutions in terms, and I think that's what we're there to do, we're there to, you know, find the problems and interpret them and present them as a risk profile to the client.” [Participant 4]

One fee-earner commented that the introduction of AI-enabled legal services reflected wider trends in society, and that while the practice of law was affected, this had not caused them to re-evaluate their own sense of self.

“Good question, not massive I suppose, because I think one feels that these sorts of improvements are happening in everyone's working life, so I think I personally consider it to be more a societal thing than a specifically legal thing. So, but at the same time, you know I understand that the practice of law now is very different in part, largely because of some of these tools and the abilities that they provide.” [Participant 16]

Further evidence that legal professionals do recognise the impact of AI-enabled legal services on the practice of law can be seen in the supplementary information from the laws of media tetrad, which details examples of how AI has both enhanced and diminished their work. Interestingly, the most frequently mentioned impacts of AI were enhancements (speed and accuracy) suggesting that legal professionals perceive AI as a traditional production technology that enhances their productivity, rather than an epistemic technology designed to generate knowledge. This might help to explain why, there was no

consistent evidence to suggest that these changes were causing a significant reappraisal of legal professionals' sense of self.

Other fee earners reported more positive feelings, indicating that technology was now an integral part of how they described their role to others, and that they felt this set them apart from their less tech-savvy colleagues and competitors.

"I guess that does mean that technology and AI has probably impacted what I think I do because I would always include it as part of a description or if not it's implicit...we are technology-led, not just technology-assisted but technology-led...I'm proud of the fact that we use that, you know, they we're at the forefront of what [Global] does in terms of technology." [Participant 18]

"I'm quite happy to tell anyone what we do and how we've used technology to make our offering better. I think you do get certain people within the firm that almost looked down on some of the stuff that we do. If that's maybe a bit controversial, but it's almost like they don't treat it like proper law. Like it's not the same way as if you were just, you know, in there with all the books and doing everything manually doing it yourself and pulling it together. I would happily have a discussion with anyone, and I wouldn't in any way feel that it's impacted my self-worth or anything like it. I'm pretty happy about the fact that I'm embracing something that's new and that I think can make my job not only easier, but help me deliver a better solution for clients." [Participant 15]

Interestingly, some legal professionals indicated that an absence of new technologies, such as AI-enabled legal services, at their firm, would be of greater concern to them as it would suggest that they were stuck focusing on administrative tasks, and not delivering the type of innovative legal services that their clients expected.

"I think it's the lack of technology that actually impacts on the identity because you sometimes you could feel like, oh, I'm, you know, I'm a part of the case, but you know that [you're not], actually thinking about the strategy, only extracting the data... If we have, for example, that software, then maybe, you know, maybe we could spend a bit more time with the supervisor discussing the strategy. And sort of we have more opportunities, maybe to

actually get to the bottom of what the legal advice is given not, you know, not just dealing with mechanical work.” [Participant 3]

“So, let's say I'm sort of a lawyer, a good one. Working at a big firm you're expected to use some of these tools...Certainly, if you were to work at a firm that didn't use it, you'd probably think yourself, probably have a negative view of yourself.” [Participant 10]

“I think probably in my fee-earning role, it wouldn't have any [impact], I would just consider that as something of we're moving with the times, and I guess I would expect as part of a really large global firm that we would do that. So I think it's just sort of integral to that role, erm. I think I would probably see smaller firms as not being as up to date on that. And I would probably think of that then, as being a bit outdated.” [Participant 20]

Participants working in non-fee earning roles, who could be regarded as producers of AI-enabled legal services were even more likely to report that their sense of professional identity, and in some instances their career development, had been enhanced. This may reflect that while fee-earners and non-fee earners both regard themselves as professionals, they conceptualise what it means to be a professional in different ways.

“In terms of my own professional identity, it's had a very, very positive impact because I never, I'm getting emotional, but I never actually really found myself fitting in one way or the other. I never wanted to practice. I was never a very good academic. I thought too wide for that. I was not good at the sort of small detail. And I wanted to be in practice. I just didn't know how. To find a role where I can actually sort of take a lot of the stuff that I was thinking about academically and actually put that in practice in terms of my own professional identity, that has been extremely rewarding for me.” [Participant 6]

“I would say [my sense of self has been] enhanced. I think I thought I was sort of joining ‘the system’ and now I see myself as a disruptor...I think I get to learn from the specialist lawyers who are also pioneers. And I think that's a pretty rare breed as most people are late adopters, so it's cool to find the early adopters.” [Participant 12]

“I think innovation was the reason I got made Partner. So, I am one of the youngest partners made in probably the last 15 years, and I am one of the quickest. So, I got my partner in eight years of qualification, which nowadays you don't hear of.” [Participant 1]

“I think one thing that a lot of people find hard when they work in this world is whether we can still identify as lawyers. Not because technology is replacing us, but because we are applying our legal skill sets to do something different. And so I don't feel like because I use a machine learning platform, it makes me less of a lawyer...I don't think I'm the same kind of lawyer that I was when I was banking lawyer, but I don't think I'm a technologist or a sales person or a non-lawyer either, because actually, having that kind of understanding of how lawyers operate and what clients expect and what's acceptable and what isn't acceptable is a huge part of my role, and I only have that if I have gone through some element of legal sense...I think you can still be a wonderful lawyer and deploy technology in your transactions, that's just you being smart.” [Participant 17]

In evaluating this theme I made the decision that the weight of empirical evidence reflected an absence of negative feelings about the impact of AI-enabled legal services on the role of legal professionals. However, the participant responses were not consistent enough overall, to justify a theme indicating that most legal professionals had positive feelings about the impact of AI-enabled legal services on their work or sense of self; although a majority of the participants in producer roles did indicate positive feelings. In relation to this theme, my own observations were not relevant to the evaluation of the theme as I have no first-hand experience of working as a legal professional.

Inferential themes

1. It is plausible to claim that legal professionals using AI-enabled legal services typically experience a hermeneutic relationship to technology.

Based on the participant experiences captured in experiential themes one and two, I inferred that other legal professionals using AI-enabled legal services would most likely describe their relationship to technology as hermeneutic in nature (Ihde, 1990). In coming to this conclusion, I noted that several participants felt comfortable in using this term

themselves in response to my questioning, rather than seeking to redefine their relationship using a different descriptor. In addition to this, the way in which participants described AI being used in the initial stages of an AI-enabled legal services to produce an output, that they were responsible for transforming into a finished product possessing meaning and value, is also consistent with a hermeneutic understanding of the relationship. While I was unable to identify empirical evidence of legal professionals experiencing hermeneutic relationships with AI in the extant literature, it has been suggested that hermeneutic relationships characterise how humans interact with AI in other contexts e.g. facial recognition technology (Hongladarom, 2020) and augmented reality (Wellner, 2020). I also felt justified in discounting the possibility that legal professionals in using the word hermeneutic, were seeking to describe a relationship characterised by machine hermeneutics (Hongladarom, 2020), in which responsibility for interpreting and assigning meaning is shared between the technology and humans. Had that been the case I would have expected participants to refer to different layers (or stages) of interpretation within the process of an AI-enabled legal services, however, there was an absence of data indicating AI playing an active role in the interpretive act.

The three acts of professional practice (diagnosis, inference making and treatment) identified by Abbott (1988) offers guidance as to why legal professionals might be keen to describe their relationship to technology as hermeneutic, while not acknowledging the related notion of machine hermeneutics. Abbott (1988) describes inference making as ‘the actual professional act’ during which the professional uses their expertise to link a diagnosis to an appropriate treatment, thus placing inference making at the centre of what it means to be a professional. In seeking to preserve their position as inference maker, legal professionals might be willing to relinquish the more mundane aspects of the act of diagnosis to AI, provided they remain responsible for interpreting and drawing inferences from the diagnostic output (Köktener and Tunçalp, 2021).

2. It is plausible to claim that legal professionals are unlikely to experience identity threat from using AI-enabled legal services.

Experiential theme three led me to infer that legal professionals are unlikely to experience identity threat from using AI-enabled legal services. The concept of identity threat refers to

experiences that are understood by individuals to bring harm to the value, meanings or enactment of an identity (Petriglieri, 2011). Of the five different categories of identity threat highlighted by Jussupow, Heinzl and Spohrer (2018) there is an absence of empirical evidence to suggest that these are being experienced by the legal professionals in my research. At the individual-level participants did not report that AI-enabled legal services were challenging their expertise or undermining their status, rather AI-enabled legal services were seen to reduce the more administrative aspects of legal work, while leaving more complex aspects of legal work intact. As a group, fee-earners in consumer roles did not report that their professional role was being significantly curtailed or redrawn in ways that were having a negative impact on their sense of autonomy, influence or professional values. In contrast, non-fee earners in producer roles indicated that their career opportunities were increasing because of AI-enabled legal services and enabling them to influence their organisations in meaningful ways. It is, therefore, my evaluation that (the absence of) identity threat is a useful concept for understanding the experiences of legal professionals using AI-enabled legal services. In relation to producers of AI-enabled legal services there is some empirical evidence of positive identity construction taking place, although the data lacks the detail necessary to infer whether this identity shift reflects a change in institutional logics or is the result of their experience of AI-enabled legal services.

Dispositional theme

The inferential themes are dependent upon the existence of a user-tool perspective of technology amongst legal professionals.

The technology as a tool perspective regards technology as a human-created tool the existence of which is to serve a human-defined purposes. Schuetz and Venkatesh (2020) have argued that these assumptions are now being challenged by AI technologies. I, therefore, decided to reflect on whether the human-tool perspective was still able to explain the empirical evidence in my own research, or whether an alternative perspective, using a different set of assumptions, was necessary to conceptualise the relationship between legal professionals and AI software within an AI-enabled legal services. For example, treating AI as an active agent, capable of acting in its own interests, rather than a passive tool (Demetis and Lee, 2018; Anthony et al, 2023).

The type of hermeneutic relationship described by legal professionals using AI-enabled legal services is firmly grounded in the user-tool perspective. The described relationship positions AI-enabled legal services as being designed by legal professionals such that AI technology is required to perform a task according to rules and parameters set by the system's designer, the output of which is then given meaning by the legal professional and used as they see fit. In many instances 'tool' was the term that the legal professionals used to describe their experience of using AI as part of an AI-enabled legal services.

"I think ultimately, it's an application, isn't it? You know, it's a tool which you use. In the same way that I'd, to take an example of the pens working, I trust that it would write things down. If the application is working, I trust that it would, you know, process the information correctly. But as I direct it to." [Participant 4]

The lack of evidence for identity threat emerging from the use of AI-enabled legal services is also consistent with AI-enabled legal services being understood through the user-tool perspective. User-tool relationships afford legal professionals a high degree of agency over the use of AI-enabled legal services, which it is assumed will produce an output that is comprehensible to the user. This type of relationship can, therefore, be seen to reduce the likelihood of identity threats emerging through challenges to professional expertise or autonomy. Had AI-enabled legal services not conformed to the familiar user-tool relationship, for example through playing an active role in core professional activities such as inference making, I anticipate the likelihood of identity threat would have increased as the status of individual legal professionals was challenged by technology capable of encroaching on activities that are central to what it means to be a professional.

It is, however, worth noting that the empirical evidence provided some tentative indications that the user-tool perspective could change in the future. As I became aware of this during the process of interviewing, I sought to explore this further by asking participants whether they thought that the output produced by AI (for legal professionals to interpret) already possessed a layer of interpretation generated by the AI. When asked to reflect on this and consider the implications of AI undertaking interpretation a mixture of responses arose. Some legal professionals maintained their conviction that only legal professionals could

provide interpretation at the moment, although they recognised this might change in the future.

“My understanding is that we've only just started using machine learning in the stuff that we're doing...And so that means that while as I say, I would describe it as a tool. It pulls out data, it presents data to us, but it's not got that [the ability to interpret]. So, when we do due diligence, you know, it ends up as a red flag; it doesn't tell us it's a red flag, or if it does tell us, it's a red flag, which it sometimes does, it's because we've coded the logic to say if that question is answered in that particularly way that's a red flag. So, we are still the ones that are interpreting...But I suppose maybe once we got it to that stage, maybe our view would change slightly. But I think at the moment where we are it's definitely still us kind of doing that analysis.” [Participant 17]

Others were already willing to accept a degree of interpretation had taken place by the time the data was presented to them, although they believed most of the interpretation was still the preserve of lawyers.

“I mean, I think there is a layer of interpretation because, for example, with the title tool it categorizes risk so it will skim the information on the basis of the Title Register and then it will pull it into different categories. So, I suppose then it just comes down to how much you rely on the technology as having done it correctly and without having to go back and check yourself or sample to see that it's done the exercise. And if you did rely on it, then I guess you could say that the very, very initial layer of interpretation has been done and you're then just adding on the advice and refining it” [Participant 19]

Hence, while the user-tool perspective is important to understanding the data generated in this research project, and by implication how most legal professionals currently understand AI, there is evidence to suggest this might change in the future. At this stage, however, the user-tool perspective remains dominant amongst legal professionals, irrespective of their specific job role (i.e. producer, consumer or liminal) or the law firm they are working for.

Dispositional theme two: A logic of protective connectedness amongst legal professionals.

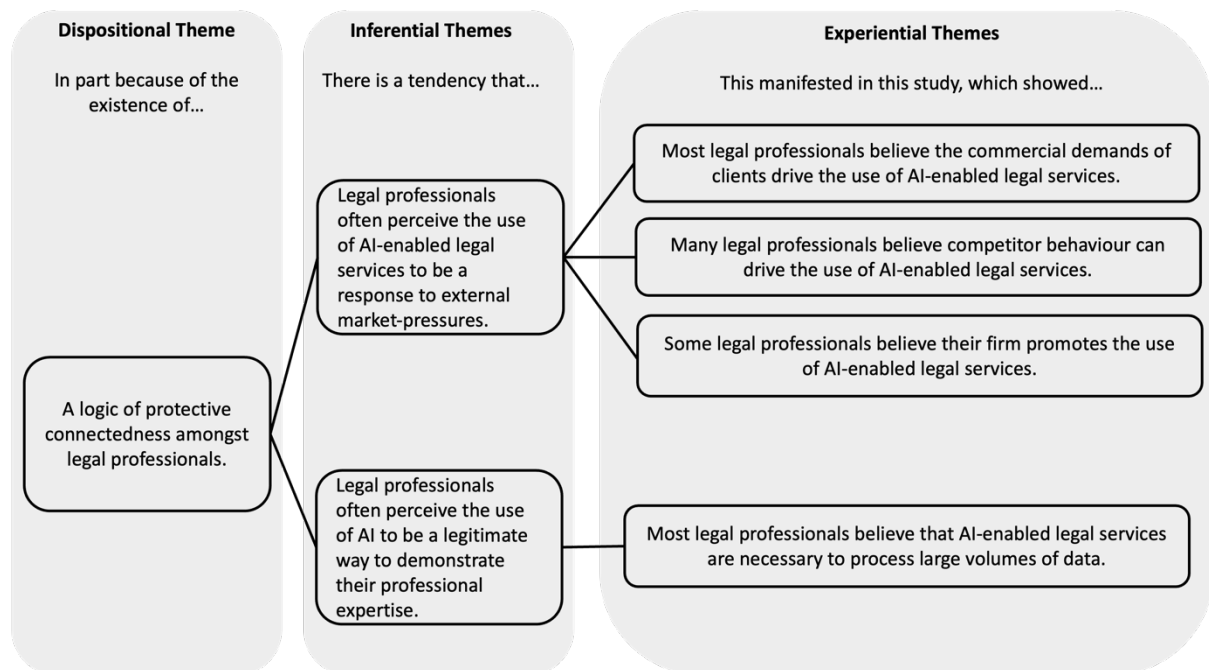


Figure 12: Explanatory statement of dispositional theme two

Experiential themes

To gain an understanding of what factors legal professionals perceived to be driving the use of AI-enabled legal services I asked them to explain how AI-enabled legal services became part of their professional practice and the goals its implementation was intended to support. Most participants in the research indicated their belief that there were multiple factors driving the use of AI-enabled legal services, with the role of clients and market competition highlighted most frequently. A minority of participants also identified the active role played by their own firm, alongside these outside pressures.

1. The data show that most legal professionals in this study believe the commercial demands of clients drive the use of AI-enabled legal services.

Most (11 of the 21) legal professionals indicated that client-related factors were significant in driving the use of AI-enabled legal services. For example, when determining which law firm to instruct, it is increasingly common for clients to explicitly ask how firms use

technology to innovate their delivery of legal services. Firms have, therefore, sought out opportunities to implement new technologies, such as AI-enabled legal services, to satisfy this aspect of client demand.

"I think there were murmurings and stuff. So [partner colleague] was getting it in the neck from the Board with clients saying, 'Well, we want firms to be innovative', (even though they didn't know what that looked like for them)...So, we have always had an eye on process and efficiency and that kind of thing and known what we've had to do to get that. So, there's always that drive from clients to reduce price and to increase service at the same time."

[Participant 1]

"Clients were just more and more expecting that we would be able to provide some sort of end-to-end solution for them...I think clients now just expect there to be tech solutions and it's almost just a kind of part of tendering. And it's an expectation now that tech will be something that we as a law firm can provide." [Participant 6]

"Everybody's talking about it [Artificial Intelligence] and, you know, the clients are getting on board...so I think that that would be a big driver for the firm, and obviously kind of monetary savings as well. But I think clients at the forefront, you know, if they like it, and if they like the sound of it, we'll do it, especially in the current market, it is so competitive."

[Participant 8]

It was also suggested that the demand for legal professionals to use new technologies, such as AI, reflects wider changes that are taking place within the client's own organisation. Hence, where clients are regular users of technology themselves, they are more likely to demand their suppliers adopt a technology-driven approach.

"Clients want to see this [Artificial Intelligence], they're doing this in their own business and they want to see people who are willing to be agile and embrace new ways of working, and they want to see lawyers who want to make the client experience better." [Participant 21]

Beyond explicit requests that law firms demonstrate the innovative use of technology, legal professionals also indicated that client demands for greater efficiency had led their firm to explore AI-enabled legal services as a potential way to achieve this. These participants

suggested that clients were primarily focused on reducing their overall legal spend, meaning they were more focused on the cost of the service, rather than how the service was delivered.

“I think from the clients' perspective, though, it doesn't really make a difference. I think they see the results and they're not necessarily bothered sometimes [where the efficiencies come from], which is fine. Although having said that, some clients, on tenders, I hear like the idea of people using cutting edge technology to get the legal results.” [Participant 4]

“I work a lot for banks and big accountancy firms, and to be honest...[they] aren't all that interested in what technology we use. What they are interested in is if we can do the job cheaper because of it. So, you know, they don't necessarily, aren't that bothered about what it means behind the scenes or what it does for us. But, as long as we can say because of that [AI] we can put less lawyers on; or the less time it takes to do your job; and this is the output you get then, yes, you know, their ears perk up, and they're very excited about it.”
[Participant 15]

At a descriptive level I felt this theme accurately captured the client-centric nature in which legal professionals practise commercial law. This was something I observed first-hand working in a law firm environment – when a client asked for something it was treated as a necessary request, not something to be questioned.

In assessing the empirical adequacy of the theme, I felt justified in attributing the theme to ‘most legal professionals in this study’, as while recognising that only just over half of participants specifically highlighted the role of clients in driving the use of AI-enabled legal services. There was also limited evidence from the interviews of dissenting voices arguing that the views of clients were insignificant, or that clients held negative views about the use of AI. The phrase ‘commercial demands of clients’ also seemed to capture the fact that clients were primarily interested in AI-enabled legal services as a means to reduce their legal spend, as opposed to seeking a change in the nature of the client-law firm relationship. I also felt warranted to describe client demands as ‘driving’ AI use, given this was the most frequently cited factor amongst the participants. The empirical evidence from the narrative timeline also highlighted several instances in which legal professionals received positive

feedback from clients about the use of AI-enabled legal services (and very few examples of clients resisting its use). Thus, further supporting the identification of clients as an important factor in explaining the widespread use of AI-enabled legal services.

2. The data show that many legal professionals in this study believe competitor behaviour can drive the use of AI-enabled legal services.

Responding to competitor behaviour to maintain a favourable position within the market for legal services, was highlighted by many (9 of the 21) legal professionals as contributing to the use of AI-enabled legal services. The use of such services was perceived to be an opportunity to create a point of difference that other firms could find difficult to imitate. Legal professionals perceived AI use to be useful both in terms of positioning the firm as forward-thinking and offering a competitively priced, high-quality service.

“Sometimes it's not necessarily about being financially competitive. It's more about building those connections...So whilst we might not be directly billing them per hour, or it's not direct and profitable, let's say; we're always first and foremost in their minds as being a firm that is able to provide lots of solutions...and therefore, when we come to tender with them, we're seen as being a tech-focused firm, or someone who is very forward thinking.” [Participant 6]

“I guess it's not wanting to be left behind. I'm sure when the quill pen started to look a little old fashioned there would be some people who were holding their ink pot jealously and thinking, ‘Well, I don't want to do it differently’. I don't want to be that person, even though I'm not by any means at the start of my career. The way to guarantee personal obsolescence is not to try and move with the times and actually adapt and develop.” [Participant 7]

“We started winning work away from other firms whose technology hadn't been as efficient as ours. And it came to [Global] because of the good work that we've been doing over that past year on those projects...When I went on secondment to [banking client]...there was a lot of chat. Not just from [banking client], but within the panel firms about what we were doing, what we could do.” [Participant 14]

It was also highlighted that some AI-enabled legal services can only be used when both parties in a transaction are willing to use them; in these instances, the use of AI-enabled

legal services requires competitor firms to collaborate with one another and agree to use the same technology to secure a positive outcome for each firm's client. The implication of this would seem to be that law firms that are unwilling (or unable) to use AI-enabled legal services are perceived negatively by clients when compared to firms that already use these services.

"It was actually the other side who suggested it and because obviously it's been incorporated within our firm, and the more senior people within the team knew that I had a lot of exposure to [named AI software], they were happy to adopt it and to use it for the transaction...So, I think all of the parties were just really happy to use that because I think everyone has seen the benefits of it." [Participant 21]

In evaluating this theme, I felt the description 'competitor behaviour' recognised that the use of AI-enabled legal services can be driven by both a desire to differentiate a firm's service offering from the wider market, and the need to collaborate with other firms through adopting shared technological solutions. As it was only a significant minority of legal professionals that highlighted competitor behaviour as a driver of AI-use, I adopted more tentative language describing competitor behaviour as something that 'can drive' the use of AI-enabled legal services.

3. The data show that some legal professionals in this study believe their firm promotes the use of AI-enabled legal services.

In addition to the above external factors, some participants (8 of the 21) in the research were keen to stress the role of their own firm in promoting the use of AI-enabled legal services. In trying to explain why the firm was motivated to do this, several different suggestions were made, with several participants highlighting different goals the firm was working towards.

Participants in both National and Global recognised that the firm's enthusiasm for AI-enabled legal services was partly explained by wider market pressures, but they also highlighted the positive experience that the increased use of technology could provide to professionals working at the firm.

“I think [AI] came about maybe four years ago with [named colleague in leadership role]. I think they just saw a real opportunity to embrace tech, and start to make lawyers’ jobs a little bit easier. [Named colleague] comes from a legal background, so she’s practiced. So I think she knew a lot of the different ways that tech could actually be used from a really practical perspective. And I think that’s kind of where it came about. They just saw the potential of it to really set us apart from our competitors. And also just to enhance the experience of the lawyers within our firm.” [Participant 6]

“So I think there’s a lot of good noises from the top of the business, that are like I guess we have to do this and it’s and it’s the way. I completely buy into that, that’s the way to future-proof the firm; clients like it; it is the way to give efficient services. It is the way to do all of these things in the right way and makes life easier for ourselves occasionally as well.” [Participant 10]

Participants from Global also suggested that their firm’s promotion of AI-enabled legal services was reflective of a wider innovation culture at the firm. This could be seen through the behaviour of senior members of the firm who actively look for new opportunities to innovate.

“I think a big one is like the culture within your firm that you’re working in because I know that [Global] are, you know, seen as quite innovative. And that was a big reason why lots of people and myself joined the firm. And I think there is a culture of trying to be innovative and embracing those things. And then that’s like a broad, overarching level. But then you’ve got to have it down on the ground as well. And I think that when I see, or talk to people, like Partners, for example, that are keen on innovation actually in practice.” [Participant 11]

“In 2016, obviously, when I started, I did my vac scheme with [Global]. There was this innovation presentation, I just remember so clearly that the Partners from other teams that had come to network with the vac schemers, you know, they were kind of sceptical and maybe afraid of this kind of word innovation and the change that could come with it... [But] by the time I joined the firm as a trainee, there was an understanding that basically we just need to improve margins and that [AI] is a great way of doing that. So that kind of fear had

turned into a sort of opportunity seeking attitude towards it, and I just found that really interesting.” [Participant 12]

Analysis of these responses by firm seems to reveal a difference in perception amongst legal professionals at National and Global. Despite both firm’s receiving external recognition for their use of technology, the role of the firm in driving the use of AI-enabled legal services was only mentioned by a minority of participants (2 of 8) at National. In contrast, almost half (6 of 13) of the participants from Global highlight the firm, potentially reflecting a stronger shared culture of innovation across the firm.

The lower number of responses on which the above analysis was based, and the differences observed between the two firms leads me to be more tentative about claiming law firms are a key driver of AI-enabled legal services. In addition, some of the participant responses explained their firm’s use of AI-enabled legal services as a response to wider market pressures; meaning the firm’s adoption of AI might reflect coercive pressure from clients or an attempt to imitate the behaviour of successful competitors (DiMaggio and Powell, 1983).

4. The data show that most legal professionals in this study believe that AI-enabled legal services are necessary to process large volumes of data.

To understand the role of AI technology within the overall process of AI-enabled legal services I asked legal professionals to explain what sort of tasks AI was used for, and the centrality of these tasks to the overall process. The size and complexity of the matters undertaken by Global and National frequently require the processing of vast quantities of client data. The ability of AI-enabled legal services to process data in ways that humans cannot easily replicate was highlighted by most of the participants (11 of 21). This has meant the use of AI is now intrinsic to certain legal tasks, either because the task is beyond the limits of human cognition, or humans alone could not undertake it in a way that would be acceptable to clients (in terms of cost and time taken).

“I’m sure, [named colleague] told you about this sort of large volume contract analysis that we sometimes get sent in. Where we have to train up a team really quickly. But it means that we can put, hundreds, thousands of documents through [named AI software], which we

previously would never have done. We would have just had to have selected however many [documents] and stuck a bunch of trainee lawyers in a room and been like, you know, you're going to be there until five a.m.” [Participant 6]

“So [named AI software] definitely it all centres around how you search for certain keywords or documents within what could be thousands and thousands, hundreds of thousands and millions even of documents. So, the AI is very central to how you go about that process... Yeah, I don't think humanly you would be able to, particularly when you get into the millions of documents, do that physically.” [Participant 9]

“Near the beginning of a case, you're overwhelmed by how much data there is, and I suppose that we use [named AI software]... [to] give you a facility where you can actually digest the information. Otherwise, it would be too impossible for someone else to do, so for like a human to do it.” [Participant 11]

“I can see the benefits of it and the sheer volume of documents that AI could deal with and, you know, make sense of, you know, splitting sort of 20,000 contracts into, those that were high risk and low risk and which ones, well, you know, foreign law and would be almost kind of group them together, and then it could like highlight particular clauses that were relevant and you know that were sort of not standard and things like that. So that was pretty impressive.” [Participant 16]

In those instances where legal professionals did not recognise AI as necessary for their work, they often explained this in terms of the bespoke nature of what they did, contrasting it to the more standardised, high-volume work of their colleagues.

“They have [used AI] because they have got thousands and thousands of essentially the same case over and over and over. They've been able to automate it. But that is not a typical experience for our team. It tends to be you get one giant case and so you need to basically make everything bespoke for that one giant case...and the bit that you could automate of that is probably quite small. So because it's quite small, it's not worth doing. And then the rest of it is the advice which we would do, well I don't think the robots are up to it quite yet!” [Participant 10]

While this example illustrates that AI-enabled legal services have not had an impact on all types of legal work, there appears to be a consensus that AI-enabled legal services offer significant benefits in those instances where volumes of data are high, and the data requires processing in some way. This means AI can impact the work of most lawyers working in firms like Global and National.

In assessing the empirical adequacy of the theme, I attributed the theme to ‘most legal professionals in this study’, as while there was a suggestion that litigation-related work might not lend itself to the use of AI (Participant 10), another participant who also worked in this practice area (Participant 11) indicated AI was of direct relevance to their work. The need to use AI-enabled legal services, therefore does not appear restricted to specific practice areas within Global and National. However, because only just over half of participants indicated the necessity to use AI, I chose to restrict its relevance to those tasks that involve ‘processing large volumes of data’, with the empirical evidence suggesting that this was likely to be high-volume work of a standardised nature. This inference was also supported by data from the tetrad exercise, where AI’s ability to allow legal professionals to work at scale was mentioned by a majority of fee-earners in consumer roles.

Inferential themes

1. It is plausible to claim that legal professionals often perceive the use of AI-enabled legal services to be a defensive response to external market-pressures.

In highlighting the behaviour of clients and competitor firms as important factors driving the use of AI, I inferred that legal professionals are sensitive to the commercial dynamics of the market for legal services, and that these external forces are perceived to play a more important role in driving the use of AI than forces within their own firm. The pre-eminence given to market pressures in the minds of legal professionals also suggested to me that managerial logics such as efficiency and profitability are highly salient in explaining the behaviour of legal professionals working in large commercial law firms. Hence changes to professional practice, such as the use of AI, are primarily a defensive response designed to protect the commercial position of the firm in terms of its market share and profitability. The more limited evidence suggesting that law firms are also willing to collaborate with one

another in their shared use of AI to treat complex cases can also be interpreted as a defensive response designed to maintain their jurisdictional boundaries and prevent complex work being undertaken by other professionals. In contrast, there is less supporting evidence to suggest the use of AI reflects a law firm's desire to improve the experience of their legal professionals. Had employee experience been a significant driver of AI, this might have been reflected through reference to the importance of the material characteristics of AI as a technology, such as 'perceived ease of use' and 'perceived usefulness' (Davis, 1989) in driving the use of AI-enabled legal services. In conclusion, the use of AI-enabled legal services amongst legal professionals is primarily a response to external forces operating at the macro-level, rather than the nature of the relationship legal professionals have with the technology.

2. It is plausible to claim that legal professionals often perceive the use of AI to be a legitimate way to demonstrate their professional expertise.

The use of AI-enabled legal services was not made mandatory in either National or Global. This suggests the widespread use of AI amongst legal professionals is more likely to reflect the outcome of a localised process of negotiation driven by professionals, than being indicative of an attempt to formalise working practices by the firm. This helps explain why legal professionals have been able to effectively coordinate their work in ways that allow them to preserve their autonomy and demonstrate their professional expertise, while working at a scale that was not previously possible. This interpretation is further supported by the evidence that indicated legal professionals were not required to introduce AI-enabled legal services to undertake bespoke work that they believed it was not well-suited to. AI-enabled legal services can, therefore, be seen as compatible with existing professional logics, rather than in conflict with them.

Dispositional theme

The inferential themes are dependent upon a logic of protective connectedness amongst legal professionals.

In seeking a plausible explanation for the two inferential themes, I recognised I would need a concept or theory that was capable of offering insights into how professionals respond to changes taking place both within and beyond the boundaries of their profession. From my knowledge of the literature relating to the professions, the concept of professional logics, which are recognised to shape the perceptions and behaviour of professionals, seemed capable of offering relevant insights.

The concept of hybrid professionalism (Noordegraaf, 2015) offers a partial explanation of the forces legal professionals understand to be driving the use of AI-enabled legal services. The accounts of participants reference the role of both professional and managerial logics. Legal professionals indicating their use of AI was a protective measure designed to ward off external financial pressures facing the profession (a professional logic); while also highlighting AI as a means of achieving greater cost efficiencies (a managerial logic). This suggests that traditional models of professionalism cannot in isolation accurately explain the experiences and behaviour of legal professionals at National and Global. However, the tension that is typically associated with trying to manage the contradictions between professional and managerial logics (Noordegraaf, 2020), was not readily apparent in the account of participants. There was, for example, no discussion of AI's ability to drive organisational efficiency being detrimental to the quality of the service being offered (a traditional concern of professionals). Rather it was highlighted that in some circumstances the use of AI permitted more sophisticated analysis of data than could be achieved by humans alone, meaning the quality of service offered to clients could be enhanced by using AI. Additionally, the introduction of AI was not necessarily seen to be accompanied by the hierarchical and more rigid working practices, which can be a feature of environments that organise work according to more managerial logics. Instead legal professionals were given a choice as to whether they wanted to adopt the use of AI in different aspects of their work, rather than being told they had to work according to a pre-defined process. I, therefore, concluded that hybrid professionalism could not provide an adequate theoretical explanation of the data generated in my research.

Going beyond hybridity, Noordegraaf's (2020) concept of connected professionalism provides a further way to try to understand the data. Connected professionalism seeks to

better explain how professionals resolve the challenge of working across traditional jurisdictional boundaries and their associated logics. Rather than seeing professional acts as something performed by professionals, they are understood to be a part of larger service processes, in which many actors, including clients, and factors, such as AI, play a role. This focus on collective processes is arguably reflected in the structure of both AI-enabled legal services and the multi-disciplinary teams that are required to design and deliver them. A shift towards connected professionalism might also help explain the willingness of participants in the research to coordinate their work with legal professionals from other functional specialisms.

Connected professionalism's focus on the relational aspect of professional work also potentially provides insight as to why the legal participants in this research identified the views of clients as playing a significant role in the introduction of AI-enabled legal services at National and Global. However, upon closer inspection the type of influence clients were having seems better described by what Bourmault and Anteby (2016) describe as 'turning inwards' to oneself, which describes professionals responding to their own direct experiences with clients, for example when a client asks for details about a firm's use of AI in tender requests. This is in contrast to what they term 'turning outwards' to the client which is characterised by professionals focusing on the actual relationship they have with clients and looking for ways to strengthen the connection between them. This type of behaviour, which if apparent would have supported the case for the behaviour of legal professionals being guided by a logic of connected professionalism, was not evident in the responses of participants. Indeed, the most notable example of legal professionals reimagining their relationships with outside parties related to collaboration in the use of AI between firms when working on opposite sides of a transaction. Such collaboration can arguably be understood as a defensive response, designed to protect the interests of both firms, which would be more characteristic of the logics of protective professionalism (that seeks to shelter the work of professionals from competitive forces), than the openness associated with connective professionalism.

This has led to alternative interpretations of the connective practices that emerge through the use of AI by legal professionals, as a defensive response to the pressures facing

professionals, what has been called ‘protective connectedness’ (Faulconbridge, Folke Henriksen and Seabrooke, 2021). The regular references amongst participants to their use of AI being a response to the behaviour of others, most frequently clients but also competitor firms, reflects the defensive nature of the adoption. This also helps explain why commercial imperatives are used to explain the use of AI. Rather than seeing AI as a way to create new services or forms of client relationships that might lead to the development of new markets, AI is primarily seen as a way to preserve the firm’s existing market share and historic client relationships. At the level of individual professionals, protective connectedness also provides a theoretical explanation as to why legal professionals are happy to make it part of their professional practice; with AI being understood as a means to ensure legal professionals (in producer, consumer and liminal roles) can continue to derive value from their existing expertise, despite it being delivered through a new set of processes.

Dispositional theme three: A technology pipeline approach to developing AI-enabled legal services.

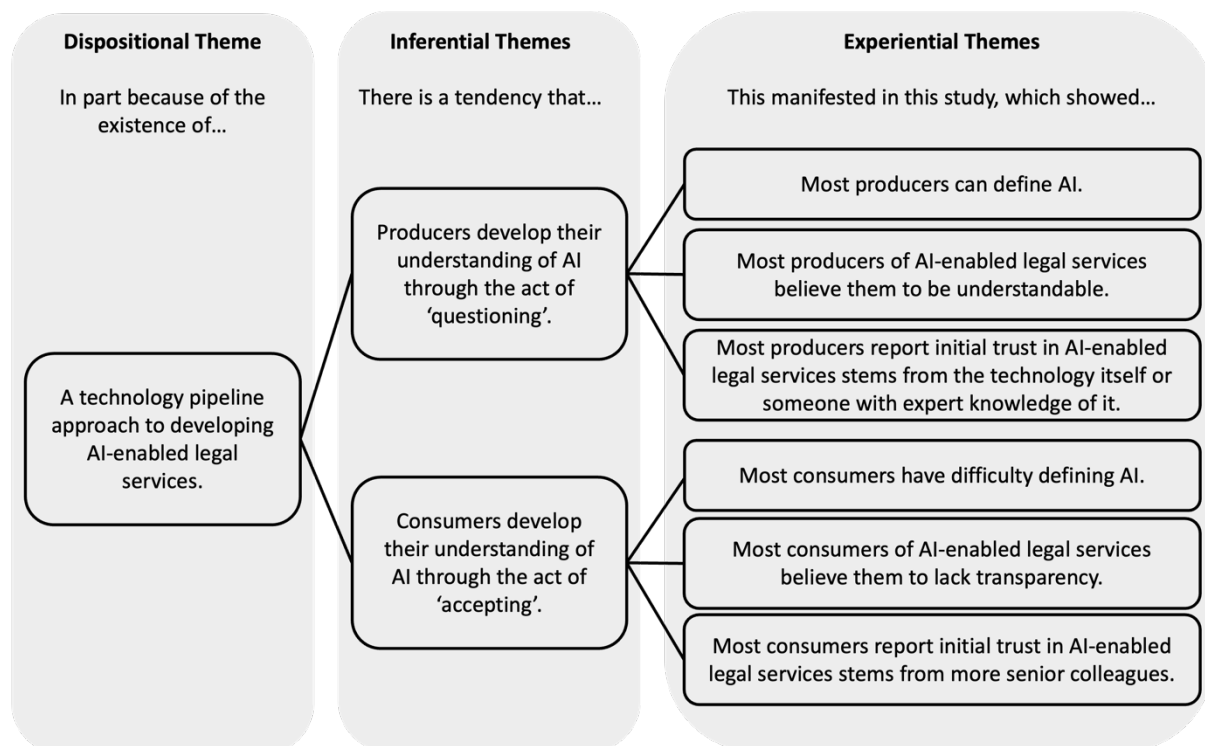


Figure 13: Explanatory statement of dispositional theme three

Experiential Themes

Before focusing on specific AI-enabled legal services, I wanted to better understand what the term Artificial Intelligence more generally meant to legal professionals. I therefore asked them to define AI in the context of their own work. This gave me a useful insight into how familiar each participant was with the phenomenon of interest.

Participants gave a range of responses that lacked any overall sense of consistency, suggesting AI may not be well understood by legal professionals. This also meant it was not possible to identify themes that reflected the views of a majority of the participants. In re-analysing the responses it became apparent that participants in producer role were more likely to give detailed answers than those in consumer and liminal roles, some of whom struggled to define AI.

Later in the interview, to gain an understanding of what role, if any, trust played in legal professionals' use of AI-enabled legal services I first asked them whether they trusted the AI technology that underpinned the services they used; I then probed further to understand how their (dis)trust in AI had developed over time.

While some participants were more hesitant than others in indicating their trust in AI-enabled legal services, they all recognised that given AI use was not mandatory at their firm, their continued use of AI indicated that they did trust it to at least some extent. In probing the reasons for their trust, a variety of responses were given, with no single reason cited by a majority of the participants. In analysing the major reasons for trusting AI I once again noticed that if the participants' responses were analysed according to whether they were a producer or consumer/liminal of AI-enabled legal services a clearer picture emerged. Most producers indicated that they believed AI to be understandable (transparent) and that their initial trust in AI was contingent on them understanding how the technology worked, or, in lieu of that, trusting the judgement of an AI subject-matter expert. In contrast most consumers and liminals indicated that AI lacked transparency and that their initial reason for trusting AI reflected their trust in the judgements of respected colleagues at their firm.

The findings of this part of the thematic analysis are presented in order to assist comparison between the groups of producers and consumers, while still showing the causal links between the different types of themes. The experiential themes of each group can be seen as ‘two sides of the same coin’ despite appearing separately in the explanatory statement. For example, the themes ‘most producers can define AI’ and ‘most consumers struggle to define AI’ arose from analysis of responses to a common question. The analysis that led to the identification of pairs of related experiential themes is presented together to aid understanding. The data relating to those participants in liminal roles is also referenced as part of the analysis. While liminals are not explicitly mentioned in the emerging themes, their responses were similar to those given by consumers, suggesting their understanding of AI was closer to that of consumers than producers.

1.a) The data show that most producers of AI-enabled legal services in this study can define AI.

1.b) The data show that most consumers of AI-enabled legal services in this study have difficulty defining AI.

In total I judged most (11) of the participants to have had significant difficulty in answering the question, this group reflected a majority of the consumers (6 of 10) and liminals (5 of 6) in the research. In contrast, none of the producers (0 of 5) displayed an inability to define AI.

Four of the 11 legal professionals struggled to differentiate AI from more general processes, such as automation, while another offered a very high-level definition that might equally have been applied to other technologies.

“I'd say it's using technology or computer systems to enhance a task by removing things that you would do. Automating things that you would do normally, but making it more efficient by the use of technology...But I basically say it's taking out certain tasks that you do anyway and letting a computer do them for you. I guess automating the process is probably the way forward.” [Participant 11]

“Erm I think I would define artificial intelligence as a piece of software, which is helping you think and you're applying some kind of information and giving it something to go on, and it gives you something out of it.” [Participant 2]

The difficulty in delineating AI from other technology was also recognised by one participant in a liminal role, who had only started to think about what AI meant in the context of their work after going on secondment from a fee-earning role to a multi-disciplinary team responsible for developing AI-enabled legal services.

“You know, it's a really is a really good question because I think my understanding of it probably isn't the best. I think I probably would have grouped all of legal tech as being AI up till now...without ever really thinking, what is AI, and what's not. I would have just seen everything as sort of innovative technology to help with the job. [Participant 20]

Contrary to the above, four long-standing members of multi-disciplinary teams, all of whom were in producer roles (and one fee-earner who had experience working in a technology start-up), offered more detailed definitions that differentiated AI from other types of technology, indicating a greater level of understanding of AI. This knowledge seemed to reflect how long they had been using AI-enabled legal services and/or their role as a producer of AI-enabled legal services. Participants in liminal roles who had only spent a short amount of time in a multi-disciplinary team, struggled in comparison.

“My understanding of it is the ability to teach the technology to anticipate and spontaneously deal with data. In a way that we don't need to continue to teach it...I see the term AI applied to, quite flat technology where we literally are just filling in an online due diligence questionnaire, which is just easier to format than a word document, and to me, that's not artificial technology, that's just tech and technological innovation that allows us to speed up our time. But ultimately, the work is being done by the lawyer, by the human because they're making an assessment on the clause that they're reading. Whereas my interpretation of AI is that the work is being done by the computer and the lawyer then reviews.” [Participant 12]

“I think for us, it's anything where we are using technology or systems to replicate human processes, so that can be anything from, I suppose that involves kind of the rules based expert systems. So, the Susskind kind of dawn of innovation right up through to the machine learning algorithms and predicting things and kind of everything in between. ...So ours tends to be around extracting data, using data to do things and then making decisions, it kind of replicates the jobs that lawyers are doing anyway” [Participant 13]

“Artificial intelligence ultimately is the aim of getting computers to apply intelligence and reasoning to process sort of more independently kind of thing rather than being told exactly what to do by the operator of the system. And obviously there's a range of artificial intelligence from AI on a specific thing like intelligent search, for example, right the way through to general AI where the idea is to have a computer, that could sort of think and behave as a human kind of thing.” [Participant 16]

“I think often people confuse artificial intelligence with machine learning. And I think machine learning is an aspect of AI. But it isn't just what AI is. To me artificial intelligence from what I've read is any computer-driven process where to a certain extent, a computer is trying to mimic a decision-making process of a human, or is behaving in a way that mirrors some sort of kind of human behaviour. I don't think that's necessarily always machine learning.” [Participant 17]

The difficulty some participants had in defining AI led me to follow-up my initial question, by asking participants whether they thought it was important for legal professionals to understand what AI was in the abstract, if they were to utilise AI-enabled legal services as part of their professional practice. Amongst the legal professionals I asked this question the majority indicated they did not think it was important for fee-earners in consumer roles to understand AI at a theoretical level, but that knowledge of specific software packages was relevant.

“I think some colleagues are perfectly able to use these AI tools without understanding how they operate in the background because, well, the colleagues I have in mind display very little knowledge of how computers work at all...it's almost a sort of app culture rather than programming culture; it's now presented to you as, if I press this button then good things

happen. OK, cool. Rather than it needing to be any more sort of intensive exercise.”

[Participant 10]

“I mean, it's not something that we're ever really asked to explain. To be honest a lot of clients rarely understand it, only on a couple of projects have we directly dealt with tech people, at clients it's more been bankers and things. So, their understanding is less than ours, and ours isn't great.” [Participant 15]

“No, I don't think they need to have a deep understanding of it, but I think they have to know either who can answer the questions if they come up or I think they need to have a basic understanding of what the system is doing that they are trying to use.” [Participant 17]

There were, however, dissenting voices that felt a lack of understanding could have negative consequences that might mean the AI-enabled legal services would not be used to their full potential.

“I think they do [need to understand AI], and I think that's probably one of the challenges with incorporating AI tools in the legal sense, because if you don't understand it, you don't really see the value in it so much.” [Participant 11]

“I think somebody needs to understand it, and I think that it can't just be left to, erm, I don't think it's something that you can pass over to a technology team, or engineers and say make this happen. I think that for it to be of any use, you need to have at least some lawyers who are interested...to really harness the benefits of it and make it something that's useful to the team.” [Participant 20]

As these themes reflect the (in)ability of legal professionals to define AI, I evaluated their empirical adequacy through assessing the high degree of variability in the empirical data – most producers defined AI in detail, while most consumers and liminals did not. In assessing the descriptive validity of the theme, I felt justified in saying that consumers found it difficult to define AI in comparison to producers as this was a judgement I felt could be made based on the precision of the language the two groups used. A stronger conclusion to have drawn would have been to say that most consumers did not understand what AI was, however, as only a minority of consumers gave responses that could be regarded as

factually incorrect (as opposed to unclear), I did not feel justified in this being reflected in the theme.

2.a) Most producers of AI-enabled legal services believe them to be understandable.

2.b) Most consumers of AI-enabled legal services believe them to lack transparency.

Most producers (4 out of 5) indicated that in the context of their work, the AI that was being used was understandable to them, meaning that they were clear how the inner working of the AI-enabled service functioned.

“[I’m] just starting to understand it, realising that it's perhaps not as complicated as you thought it was, and that actually something like [named AI software], whilst the output looks very complex, once you start to sort of look behind the scenes you realise how conditional logic works. And how you can build conditional logic, and when you can start to do it yourself and you go, OK, I understand this, this is actually very understandable for me. And whilst at first it was quite an intimidating prospect to understand what it is that this tool does, actually, in reality, I understand it and I like it kind of enhances my trust in it. I think explainability and understandability is key, isn't it? [Participant 6]

“My [multi-disciplinary team] are a team where we only deploy something if we understand how it works and we can explain it... I understand why the system arrived at a particular answer, and I can unpick that and I can trace that back. So, logic-based system, or document automation, or the systems where I've given it a set of rules, I can undo these rules, so I always know why it's done something.” [Participant 17]

“Transparency is very important in these scenarios because they may impact the decision taken by a fee-earner, which then has consideration of our insurance position for an insurer. So yes, it is very important to be transparent.” [Participant 13]

In evaluating this theme, I chose ‘understandable’ rather than ‘transparent’ to describe producers’ beliefs about AI as while both words were used, ‘understandable’ seemed to also capture why transparency was highlighted as being important i.e. without transparency, understanding is impossible. The consistency with which this view was reported by

producers meant 'most' seemed to accurately reflect the groups agreement about AI's understandability.

Most consumers (6 out of 10) indicated they did not find it easy to understand the inner working of the AI that underpinned the AI-enabled services that they used. This was because in their view the AI lacked transparency. However, this did not seem to be of significant concern to those participants who indicated they did not feel an understanding of AI was part of their role, or necessary in order for them to use AI in their work.

"I don't know what technical process algorithms or systems it uses. I wouldn't profess to know in in any kind of detail how it comes to that...I'm not really worried about the sort of the process or how it gets there. I don't need that transparency." [Participant 7]

"Erm I think some colleagues are perfectly able to use these AI tools without understanding how they operate in the background because, well, the colleagues I have in mind display very little knowledge of how computers work at all...So, yeah, so I don't feel like people have to have this great knowledge of computers to be able to use the AI now." [Participant 10]

"I can't say, I have the opportunity to look under the hood of that really, that's not kind of really in my role. I do understand what it does kind of at a high level and could probably talk about it at a high level. But in terms of the specifics, probably not really." [Participant 14]

Of the six liminals in the research, four of the six felt AI lacked transparency. Interestingly, one participant in a liminal role at National suggested that the reason why consumers perceived AI as lacking transparency was because of a lack of effort on their part, rather than a flaw in the technology itself.

"I think at the moment, those who are interested and willing can certainly see the transparency...Most of the time, we use this technology for things that are quite obscure to do, and no one really wants to do in the first place. So when we look at, you know, those tasks. No one really wants to do them in the first place. And people just don't really care 'why?'...I think everything is transparent at this point. But just some people are not really taking advantage of this, I suppose, because they don't need to. And that's fair enough, [Participant 8]

In contrast a producer at Global took a much less dismissive view of their colleagues, suggesting that most consumers would be motivated to understand AI out of a sense of natural curiosity. This view was not supported by the data, suggesting producers may hold false assumptions about consumers' beliefs about AI.

"I don't think you'll find many lawyers that do trust the black box just because I don't think it's in our nature just to take things at face value. I think we're always going to want to know why something is the way that it is, because that's just the nature of the sort of people that we tend to be. And I don't think that's necessarily being risk averse, this is being inquisitive and wanting to understand how something pieces together." [Participant 17]

In evaluating this theme, I chose 'lacks transparency' to describe consumers' beliefs about AI as this seemed to be a relatively neutral way to capture what was common across most of the participants' responses. While some participants also sought to explain the reasons for this, there was not the consistency of data to conclude whether the perceived lack of transparency stemmed from the design of the technology or the limitations of the consumer's knowledge of AI (or both).

3.a) Most producers report initial trust in AI-enabled legal services stems from the technology itself or someone with expert knowledge of it.

3.b) Most consumers report initial trust in AI-enabled legal services stems from more senior colleagues.

Most producers (4 out of 5) indicated that their initial trust in AI (i.e. the trust they had prior to it entering widespread use at the firm, reflected their technical understanding of how the AI worked as a technology. This type of trust, therefore, seemed to naturally follow from their belief that AI was understandable.

"Again, I think it really depends on what sort of artificial intelligence we're talking about. So if we're talking about, something where I've told the machine to go and retrieve something or I've told it to make a particular decision, so we know with logic-based systems like [names AI system] I've told it, that if Question A is One, and Question B is Two, then you need to tell me that it's Three, erm I completely trust that. And it doesn't require any interpretation

because I have set the logic rules, and I understand that it's always going to do that no matter what and there's no chance of variation.” [Participant 17]

“Yes, I do trust it because we're, you know, we're making sure that we, we think with that transparency right from the start. And trust probably is the right word because you need to establish trust from the get-go when you're doing these projects and make sure that the results that are coming out of that everybody understands and agrees with.... So essentially at every step along the way, starting with the very basics and just building a little bit on top of that is kind of the way that I've used to demonstrate both, you know, gain trust and demonstrate they trust that you know what's going on in that process.” [Participant 13]

Most producer participants (3 out of 5) also recognised that it was not always realistic to try to develop a detailed understanding of every AI technology the firm might use, even if it were technically feasible to do so. In such circumstances, producers indicated that their initial trust in AI would reflect their trust in the views of individuals who could be regarded as subject matter experts on the AI.

“I suppose the honest answer is that I just trusted the colleagues around me when I was sort of told that [named AI system] was great and I'd never used it before. I knew about as much about it as a website could possibly tell me or a podcast that I'd listened to. And then I sort of just trusted the opinion of the specialists around me and people working for years alongside it. So I think to that extent, the attitude of the team around you is very important.” [Participant 6]

In answering this question, the same producer also flagged an important difference between their trust in the views of an expert (i.e. trust arising from expert knowledge) and how consumers were likely to develop trust, based on the seniority of the person who trusted the AI system (i.e. trust arising from their relative position to one another).

“But in terms of the teams around you or lawyers sort of adopting it, they are heavily influenced by the attitude of their superiors. Law is a very hierarchal sort of industry world, more so, I think, than a lot of other industries.” [Participant 6]

Another producer was also quick to recognise their own influence (as a subject-matter expert) on the development of trust amongst their colleagues.

“I see lawyers using [AI] who don't understand it. I think that comes from their trust in us and also seeing the results and saying, Yeah, that looks right. That doesn't look wrong. I don't know how you got that, but it looks right. And [Participant 16] says it works. And if [Participant 16's boss] says it works, it works right?”

In evaluating this theme, I sought to capture the technology-focused nature of producers' initial trust in technology. Their responses indicated that fundamentally it was trust in the functionality and performance of the AI technology that was of primary importance to them, but there was an acknowledgement that this could be evaluated either first-hand or by someone else with expert knowledge of the technology that was known to the producer. I therefore chose to reflect both of these technology-focused sources of trust in the theme.

Most consumers (6 out of 10) when talking about the development of trust indicated that their tendency to trust AI prior to using it was based on the relationship they had with colleagues, and what those colleagues believed about AI. In explaining which colleagues' views would be perceived as most trustworthy, participants flagged seniority (which in law firms typically reflects the number of years of post-qualification experience a lawyer has), or technical expertise with AI, as the major factors they would consider.

“If someone senior has seen some documents or some advice then I usually wouldn't really question that because at the end of the day, you know, it's sort of up to senior people... because ultimately they are responsible for the file. So, I think I would just trust them if they said, Oh, just follow the software.” [Participant 3]

“You know, if somebody like [senior colleague in producer role] said to me, you know, that is good, this works you can trust that I would definitely take that at face value. We use it, but we don't procure it. You know, I don't go to Legal Geek or keep up with what the market's doing. So I think if somebody said that it was good, then you would [trust it]. And I think as well, we are probably quite strict in terms of the technology we roll out at the firm, and I

would trust that if we put something into general circulation, that it's been properly tested, so I think that would give me comfort too.” [Participant 18]

In evaluating this theme, I focused on ensuring the relational nature of consumers' initial trust was captured. Their responses indicated that rather than making a direct judgement about the trustworthiness of the technology, which it might be assumed they would not feel capable of doing, they used the trustworthiness of their colleagues as a proxy for trust in AI.

Inferential themes

1. It is plausible to claim that producers develop their understanding of AI through the act of questioning.
2. It is plausible to claim that consumers develop their understanding of AI through the act of accepting.

I generated these two inferential themes by abstracting from the experiential themes using the related concepts of 'questioning' and 'accepting' (Anthony, 2018). These concepts describe two different sets of practices that knowledge workers have been demonstrated to use empirically, to determine the trustworthiness of the technology they use.

Questioning practices characterise behaviours which seek to examine the assumptions that are embedded within a technology (Anthony, 2018). The data generated in interviews with producers indicated that they had developed a relatively sophisticated understanding of AI, within the context of their own work, which allowed them to talk knowledgeably about the technology and how it worked. The data from the narrative timelines generated by these participants also highlighted that a majority of these individuals had identified 'working with external providers of AI' and 'delivering AI projects for their firm' as highly significant to their development of a positive attitude towards AI. These types of experience are likely to have afforded producers the necessary time and space to engage in questioning behaviours. When procuring technology from external providers, producers would be responsible for assessing whether the technology met their firm's requirements, which would necessitate them questioning software developers about the functionality of the software. Producers also play a project management role when rolling out new AI software at their own firm.

This requires them to engage in questioning behaviours through their involvement in both developing and testing the algorithmic models that underpin the AI software used at their firm.

Producers also indicated their belief that AI needed to be transparent and understandable, and that the AI-enabled legal services they utilised met this standard. Questioning practices are a plausible mechanism through which producers would have been able to gain the necessary knowledge and experience to unpack the AI software and make it understandable. Finally, when talking about the development of trust in AI, producers highlighted a desire to understand the technology as a means to develop trust in it. This can be understood as a layperson's description of Anthony's (2018) questioning practices, as the producer is indicating that they need to understand the assumptions of a technology, in contrast to simply treating the output of a technology as legitimate, which would be characteristic of accepting behaviours. In conclusion, based on both the actions of producers and the outcomes associated with their actions, I felt warranted in describing the activities that producers engaged in to develop their understanding of AI as 'questioning'.

In contrast to the above, accepting practices describe professionals who use technologies and trust their outputs, despite not exploring the assumptions embedded within the technology i.e. the technology remains inaccessible and 'black boxed'. The interview data generated with consumers of AI-enabled legal services indicated that most of them believe AI lacks transparency (meaning they treated it a black box). In addition to this most of them had difficulty in defining AI. The data from the narrative timelines of consumers helps explain why consumers did not engage in the same questioning behaviours as producers. While indicating that the encouragement they received from their immediate colleagues was significant in them developing a positive attitude towards AI, consumers did not reference experiences that would necessarily have offered the opportunity to engage in questioning practices e.g. collaboration with AI providers or involvement in AI projects. In spite of the above, producers indicated the development of trust in AI through engagement with colleagues whose own beliefs about AI were important in shaping the beliefs of consumers. The expertise and seniority of these colleagues were both highlighted by several consumers as important in the development of trust. This too points to consumers

engaging in accepting practices, as empirical research (Anthony, 2018) demonstrates that status differences reduce the propensity of more junior professionals to engage in questioning behaviour and make accepting behaviour more likely. I, therefore, felt justified in describing the behaviour of consumers that led them to use and trust AI as 'accepting'.

The data relating to participants in liminal roles was also suggestive of them having engaged in accepting practices as, like consumers, the majority of liminals had difficulty defining AI and indicated that they believed AI to lack transparency; yet in spite of this they used the technology in their work. However, the idiosyncratic role played by liminals at both Global and National meant their individual experiences varied from one another, and that themes in the data were harder to identify. In contrast, the experiences of consumers were relatively consistent. This made it more difficult to abstract from the empirical data to generate inferential and dispositional themes capable of explaining the liminals' behaviour. However, it was notable that all liminals indicated that the 'use of AI was encouraged by their wider team' and that this played an important role in them developing a positive attitude towards AI. In contrast, only two liminals highlighted 'working with external providers of AI', while one liminal highlighted their involvement in 'delivering AI projects for their firm'. Hence, I concluded that the experience of liminals was more similar to consumers than producers, and that the empirical evidence was more suggestive of accepting practices amongst liminals.

Dispositional theme

The two inferential themes are dependent upon producers and consumers interacting with Artificial Intelligence at different stages of the technology pipeline.

In seeking to explain why producers engaged in questioning practices and consumers engaged in accepting practices it seems sensible to consider whether the demands of their respective roles could provide an explanation, as both status and task allocation have been shown to influence these behaviours (Anthony, 2018; 2021). In explaining their reasons for trusting a colleague's judgements about AI, several consumers highlighted the AI-related expertise that their more senior colleagues had, whether that was from personal experience of using it as a consumer or having developed/tested it as a producer. There is, therefore, a

temporal element suggesting that consumers understand themselves as interacting with AI much later than other colleagues are required to. The idea of ‘technology pipelines’ being used by law firms to develop and deploy AI-enabled legal services is instructive in this regard (Armour and Sako, 2020; Armour, Parnham and Sako, 2022). The role of the producer sees them involved at multiple stages of the pipeline, during which they have several interactions with the technology. Key responsibilities of producers include design and procurement of AI software and then the subsequent testing of this software to see if it works. These responsibilities require producers to develop a detailed working knowledge of the software, hence the importance attached to transparency.

In contrast consumers typically interact with the technology pipeline in its final stage (application of results), by which point the AI technology has already been implemented and approved for wider use by the firm. It, therefore, makes sense that the initial trust of consumers (and their decision to use the technology) will be based on human-centric factors, specifically their interactions with those humans who have been involved in the earlier stages of the technology pipeline. Furthermore, engagement in questioning behaviour at this stage of the technology pipeline, could potentially be seen as inefficient and unnecessary in order to meet the demands of the consumer role. Given, an important logic for the use of AI-enabled legal services is organisational efficiency, it seems likely that the structure of the pipeline and the associated allocation of tasks between producers and consumers, will encourage producers to question and consumers to accept. It is my judgement, therefore, that the empirical differences in behaviour seen between consumers and producers reflect their differing levels of exposure to AI across the stages of the technology pipeline and the unique demands of their respective roles.

6.3 Development of technological frame for AI-enabled legal services

In seeking to answer, *How are AI-enabled legal services understood by the legal professionals that use them?* my research sought to identify the assumptions, knowledge and expectations that legal professionals held towards AI-enabled legal services. The thematic analysis presented above linked the participants’ own experiences and

interpretations of AI-enabled legal services to concepts and theories from the extant literature, with the aim of showing how these experiences and interpretations had arisen. The process of thematic analysis required each experiential theme to be presented in its own right, and while links between different experiential themes were made through the generation of inferential and dispositional themes, the analysis meant each finding within the data was considered in turn, rather than simultaneously. In answering the over-arching research question there was, therefore, a need to bring these findings together and present them in a more holistic way.

The empirical data generated by each participant can be understood to reflect their own mental model of AI-enabled legal services, from now on referred to as their technological frame (Orlikowski and Gash, 1994). Hence, the patterns in the data identified through the thematic analysis, reflect areas of similarity (and difference) in the frames of the legal professionals in this research. The frame structure indicates what factors relating to AI demanded the individual's attention during the process of interpretation, with the content within this structure reflecting the individual's understanding of AI. Furthermore, as technological frames have been demonstrated to endure over time, when the context within which they developed remains unchanged (Davidson, 2002), there is value in understanding the structure and content of frames, as they can play an important role in shaping the future behaviour of legal professionals towards AI-enabled legal services.

The technological frames used by legal professionals to make sense of AI-enabled legal services are presented below using an adaptation of the five-dimension framework identified by Spieth *et al.* (2021). The original five dimensions reflected those most commonly found in the extant literature; this meant they offered a useful framework through which to initially organise the participant data. Adopting this approach allowed me to identify both how the structure of the frames of legal professionals were similar to those found in other contexts; and points of difference that were novel to the context of the UK legal sector. Organising the data in this way also made it easier to see how the content of the frame was spread across different dimensions, with the dimensions that contained significant volumes of data likely to help explain the behaviour of legal professionals towards AI. Finally, where differences in the structure and content of frames were found

amongst groups of legal professionals (most notably between producers and consumers) a comparison of the technological frames of each group was made. These areas of incongruence hold particular interest as these can lead to different behavioural responses towards a given technology (Orlikowski and Gash, 1994).

Evidence of congruence in the technological frames of legal professionals

While the data relating to each participant reflected their own unique context and experiences, the technological frames of the legal professionals in this research shared similarities in both their structure and content. This suggests the participants’ interpretations of AI were shaped by a common set of factors. Of the five dimensions identified by Spieth *et al.* (2021), four showed high levels of homogeneity in their content. Two of these were rooted in legal professionals as individuals: their *personal attitude to AI* and their assessment of the *value of AI*; the remaining two reflected aspects of the wider environmental context: *firm influence* and *legal sector influence* (see Table 23).

Table 23: Shared structure and content of the technological frames of legal professionals relating to AI-enabled legal services.

Dimension	Participant understanding of AI-enabled legal services
Personal attitude to AI	<ul style="list-style-type: none"> • AI is there to augment the role of legal professionals. • AI plays a diagnostic role, legal professionals undertake interpretation. • AI is understood to be a positive for legal professionals. • AI is expected to behave like a tool.
Value of AI	<ul style="list-style-type: none"> • AI improves the speed and accuracy of legal work. • AI is necessary to undertake certain aspects of high-volume legal work.
Firm influence	<ul style="list-style-type: none"> • Colleagues are supportive of the use of AI-enabled legal services. • The firm’s culture encourages the use of AI-enabled legal services. • AI is important to the ongoing success of the firm.
Legal sector influence	<ul style="list-style-type: none"> • Clients expect legal professionals to use AI to deliver cost-effective services. • Law firms use AI to compete effectively in the market for legal services. • Using AI can help protect the roles and status of legal professionals.

The *personal attitude to AI* of legal professionals was positive in nature, with AI understood to behave as a tool designed to augment (but not replace) the role of legal professionals. More specifically, AI was expected to assist legal professionals with the diagnostic aspect of their work, while leaving interpretive tasks to legal professionals. Linked to this, legal professionals perceived the *value of AI* in terms of its positive impact on the speed and accuracy of their work. This meant that for certain types of tasks the use of AI was now seen as necessary. Looking beyond the material features of the technology, *firm influence* shaped interpretations of AI, with both the law firm culture and colleague behaviour seen to encourage the use of AI-enabled legal services. Alongside this, *legal sector influence* (seen through the behaviour of clients and competitor firms) was also found to be shaping interpretations of AI positive ways.

Evidence of incongruence in the technological frames of legal professionals

Alongside the evidence of frame congruence identified above, the findings of the thematic analysis also highlighted differences between groups of legal professionals, suggesting potential areas of frame incongruence might exist. The differences that were found between groups of producers and consumers are reflected in the content of their frames, rather than their overall structure. Specifically, content relating to the dimension *personal attitude to AI* and a further dimension of context called *expert influence*. The dimension of *expert influence* does not appear in the original framework proposed by Speith *et al.* (2021), although it shares some similarities with the dimension they name *supervisor influence*. *Supervisor influence* describes the way in which an individual's relationship with their immediate supervisor can shape their understanding of a technology. In contrast, *expert influence* reflects the relative expertise of individuals, with the views of senior colleagues (who are perceived to have greater levels of legal expertise and/or expertise with AI) affecting the interpretations legal professionals make about AI. Hence, while supervisor influence can be understood to operate within formalised hierarchical relationships that involves a superior and subordinate, expert influence operates through recognised differences in expertise amongst colleagues, without there needing to be a formal relationship between them. The identified differences in the content of the technological

frames of producers and consumers are presented alongside one another to aid comparison in Table 24.

Table 24: Contrasting structure and content of the technological frames of producers and consumers relating to AI-enabled legal services.

Dimension	Producer understanding of AI-enabled legal services	Consumer understanding of AI-enabled legal services
Personal attitude to AI	<ul style="list-style-type: none"> • AI is something I have knowledge and understanding of. • The way in which AI works is understandable. • AI can be trusted if it can be understood. 	<ul style="list-style-type: none"> • AI is something I do not understand well. • The way in which AI works lacks transparency.
Expert influence	<ul style="list-style-type: none"> • AI is understood to be more trustworthy when technical experts endorse its use. 	<ul style="list-style-type: none"> • AI is understood to be more trustworthy when senior legal experts endorse its use.

The *personal attitude to AI* of producers reflected their beliefs that AI was understandable and that this made it trustworthy. In contrast, consumers treated AI as lacking transparency, which made it difficult for them to understand. *Expert Influence* was seen to shape producer and consumer interpretations of AI’s trustworthiness. In the case of producers, the endorsement of AI by individuals with technical expertise relating to AI was salient to their decision to trust. In contrast, consumers’ assessments of trustworthiness were found to be influenced by endorsements of the technology from colleagues with legal expertise.

Drawing on dispositional theme three of the thematic analysis, which was generated to explain the observed differences in the data, the incongruence in the content of the frames of the producer and consumer groups is likely to reflect the distinctive role played by each group within the technology pipeline. The data for individuals in liminal roles indicates the content of their technological frame with respect to *expert influence*, more closely

resembled the frames of consumers, despite their role and responsibilities being more extensive, and including aspects of the producer role. The empirical data did not, therefore, suggest that liminals had a distinct technological frame of their own.

Conclusion

In seeking to answer the question, *How are AI-enabled legal services understood by the legal professionals that use them?* the overall structure and content of the technological frames of legal professionals display a high-degree of congruence. Amongst legal professionals working in different law firms and job roles, five common dimensions were identified as influencing their interpretations of AI. Four of the five dimensions also demonstrated homogeneity in aspects of their content, which helps explain why these individuals displayed similar behaviour towards AI-enabled legal services and had all chosen to make AI part of their professional practice. However, there was also evidence of frame incongruence in relation to the content of the dimensions of *personal attitude to AI* and *expert influence*, when comparing legal professionals in producer and consumer roles. These dimensions may, therefore, help to explain differences in the behaviour of producers and consumers towards AI-enabled legal services.

Taken together, the findings suggest that the interpretations of AI made by legal professionals are being shaped by a shared framework of five common factors, but that the specific interpretations that are made differ in important respects, depending on the role the legal professional performs within the technology pipeline.

7. DATA ANALYSIS – PROCESS TRACING

This chapter presents the findings of the theory-building process tracing that I undertook following the completion of the thematic analysis, discussed in the previous chapter. The process tracing focused on the development of a plausible theoretical mechanism that was capable of explaining the use of AI-enabled legal services by legal professionals at National and Global. The structure of the chapter follows the three-stage process developed by Beach and Pedersen (2016). The analysis begins with an explanation of how the two typical cases used in the process tracing analysis were selected (Section 7.1).

Having selected the typical cases, a descriptive narrative was developed for each of the two cases using the empirical data each of the case participants generated in their interviews and written exercises. The first typical case describes a producer of AI-enabled legal services; the second typical case is of a consumer. Where relevant insights drawn from the thematic analysis were incorporated into the case narrative, to complement the empirical data. Each narrative was then analysed with a view to identifying empirical fingerprints in the data that would provide support for the existence of a specific causal mechanism.

Section 7.2 focuses on the typical case of a legal professional in a producer role. It contains the case narrative, proposed causal mechanism and an evaluation of its explanatory power of the use of AI-enabled legal services. Section 7.3 provides a similar analysis for the typical case of a legal professional in a producer role.

7.1 Typical case selection

The first stage of theory-building process tracing required me to identify a typical case from a bounded set of cases, each of which was seen to contain a similar set of plausible causal factors linked to the use of AI-enabled services by legal professionals.

Table 25 summarises which of the potential causal factors was identified in relation to each of the 18 cases considered in the data analysis process. As the data show, no single causal factor was found to be present in all 19 of the cases. Two causal factors *power of AI to enhance legal work* and *power of team* displayed a high degree of coverage, but falling short of being a necessary factor for use of AI-enabled legal services.

It is important to note that these two causal factors were operationalised by combining data that reflected different dimensions of the named power, but which could arguably have been considered distinctive causal powers in their own right. In the case of *power of AI to enhance legal work* (identified in 17 of 18 cases) this combined the dimensions of *accuracy*, *speed*, *quality*, and *working at scale*. The distribution of these dimensions of powers varied across the cases, with *speed* showing the broadest coverage (13 cases) and *working at scale* the lowest (7 cases). The causal factor *power of team* (identified in 16 of 18 cases) was operationalised by combining all data that referenced the normative influence of teams on individual legal professionals; because of the way the data was coded, it was not possible to differentiate between different types of team, which would have provided a more fine-grained analysis. Had either of these factors been identified as a necessary factor for AI use, there would have been a need to evaluate whether these factors needed to be revised and made more precise, in order to maximise their explanatory power.

Table 25: Summary of causal factors relating to each case

Case	Potential causal factors									Outcome
	Power of AI to enhance legal work	Power of AI to generate trust	Power of project work with AI	Power of personal knowledge	Power of team	Power of firm	Power of professional network	Power of clients	Power of competitors	Use of AI-enabled legal services
1	+		+	+				+	+	+
2	+				+	+		+	+	+
3	+				+		+			+
4	+			+	+		+	+		+
6	+	+	+	+	+	+		+	+	+
8	+		+	+	+			+	+	+
9							+	+		+
10	+				+	+			+	+
11	+				+	+		+		+
12	+	+	+	+	+	+			+	+
14	+				+	+	+		+	+
15	+				+	+		+		+
16	+			+	+					+
17	+	+	+	+	+					+
18	+	+		+	+			+		+
19	+		+		+			+		+
20	+		+		+	+	+	+	+	+
21	+		+		+		+	+		+

The analysis then progressed to identifying bounded populations of cases within the overall population, which shared common causal factors. Table 26 summarises the similarity of each case to the Group A set of causal factors, which I selected to reflect the causal powers relating to both humans and technology in their interactions with one another, as identified in the literature relating to trust development in human-technology interactions (Hancock et al, 2011). Because the relationship that develops between legal professionals and AI was a focus of my research it made sense to consider both human and technological causes together, rather than looking at each one in isolation. The technology-related causal powers, *power of AI to generate trust* and *power of AI to enhance legal work*, reflect the *attributes* and *performance* of AI within an AI-enabled legal service. The human-related causes powers, *project work with AI* and *power of personal knowledge*, reflect the *characteristics* and *abilities* of the legal professionals themselves. Cases are ordered from most similar to least similar. Some cases such as 6, 12 and 17 showed similarity across all four factors, with a further three cases 1, 8 and 18 showing similarity across three of the four factors. In contrast, case 9 showed similarity with none of the four factors and a further seven cases showed similarity in only one area. The data, therefore, highlights different combinations of causes across the population of cases, suggesting a causally heterogeneous population.

Table 26: Summary of Group A causal factors (technology and human) relating to cases

Case	Potential causal factors				Similarity Number of shared causes	Description of case
	Power of AI to enhance legal work	Power of AI to generate trust	Power of personal knowledge	Power of project work with AI		
6	+	+	+	+	4	Producer
12	+	+	+	+	4	Producer
17	+	+	+	+	4	Producer
1	+		+	+	3	Producer
8	+		+	+	3	Liminal
18	+	+	+		3	Consumer
19	+			+	2	Liminal
20	+			+	2	Liminal
21	+			+	2	Liminal
4	+		+		2	Consumer
16	+		+		2	Consumer
2	+				1	Liminal
3	+				1	Consumer
10	+				1	Consumer
11	+				1	Consumer
14	+				1	Consumer
15	+				1	Consumer
9					0	Consumer

Table 27 summarises the individual case data for Group B, which reflected the causal powers relating to aspects of the environment that can impact trust development in human-technology interactions (Hancock et al, 2011). These potential causal powers are reflective of the broader social structure in which AI-enabled legal services exist, and so it made sense to consider these separately from the Group A causes. The causal powers, *power of team* and *power of professional network* reflect the *collaboration* dimension of the environment factor; whereas the *powers of the firm, competitors and clients*, reflect the *institutional* dimension of the environment. Cases are ordered from most similar to least similar. Only one case (20) showed similarity across all five causal factors, however, cases 2, 6, and 14 showed similarity across four of the five factors. In contrast, cases 16 and 17 showed similarity with only one of the five factors, and a further five cases showed similarity in only two of the five areas. The data demonstrates a causally heterogeneous population, but unlike Group A no particular combinations of factors within Group B is suggestive of a bounded population.

Table 27: Summary of Group B causal factors (environment) relating to cases

Case	Potential causal factors					Similarity Number of shared causes	Description of case
	Power of team	Power of professional network	Power of clients	Power of competitors	Power of firm		
20	+	+	+	+	+	5	Liminal
2	+		+	+	+	4	Liminal
6	+		+	+	+	4	Producer
14	+	+		+	+	4	Consumer
4	+	+	+			3	Consumer
10	+			+	+	3	Consumer
11	+		+		+	3	Consumer
15	+		+		+	3	Consumer
8	+		+	+		3	Liminal
21	+	+	+			3	Liminal
12	+			+	+	3	Producer
3	+	+				2	Consumer
9		+	+			2	Consumer
18	+	+	+			2	Consumer
19	+		+			2	Liminal
1			+	+		2	Producer
16	+					1	Consumer
17	+					1	Producer

On reviewing the data, I identified a tentative explanation for the distribution of causes across the cases in Group A, through relating the combination of causes to the type of interaction (consumer, producer or liminal) that each case had with AI-enabled legal services at their firm. This explanation for the use of AI-enabled legal services was also identified in Dispositional theme three from the thematic analysis – the existence of a

technology pipeline approach to developing AI-enabled legal services – which reflected different patterns of data amongst consumers and producers

When overlaying the type of interaction against the causal factor data for Group A two bounded populations of cases emerged. All *producers* (cases 1, 6, 12 and 17) were highly similar, displaying the presence of a majority of the potential causal powers relating to the human and technology factors that underpin the development of trust in technology. In contrast, most *consumers* (cases 3, 9, 10, 11, 14 and 15,) were highly dissimilar to the producer group (but highly similar to one another) with the majority of the potential causal powers absent in these cases. The liminal cases could not be considered to comprise a bounded population of their own as they were less similar to one another as a group, with some liminal cases more similar to the bounded population of producers and others more similar to consumers.

In identifying typical cases for the two bounded populations in Group A I focused first on the degree of similarity shown by the case to its bounded population. This led to the initial selection of case 17 as most typical of the cases within the producer group. I then made an assessment as to whether there was enough rich empirical data contained within the interview and written exercises of the case to generate a case narrative. I judged case 17 to contain sufficient data from which to develop a case narrative.

Amongst the group of cases which showed a high level of dissimilarity from the producer group, case 9 was most dissimilar, but I also judged it be lacking in the empirical richness required to develop a case narrative. The remaining cases in the consumer group all displayed a similar level of dissimilarity to the producer group and so based upon the empirical data available for each case, I identified case 14 as offering the richest data from which to develop a case narrative from.

In contrast to Group A, when overlaying the type of interaction (consumer, producer or liminal) against the causal factor data for Group B no clearly bounded populations emerged. This suggested different combinations of environmental causal factors were of potential relevance to each case, and that specific combinations of causal factors were not associated with the position of the individual within the technology pipeline. As no bounded

population of cases emerged, it was not possible to identify a typical case from which to develop a case narrative. Further explanations that might explain the distribution of cases in Group B, were sought, but none led to the identification of bounded populations within the data.

The second stage of theory-building process tracing involved the creation of a case narrative for each of the typical cases based on the empirical data related to it. The available data comprised the content of the interview and written exercises completed by the case participant. The aim of this step of the analysis was to identify empirical fingerprints that could be inferred from the empirical data in the narrative. These fingerprints would then be used to infer the existence of a causal mechanism(s) capable of linking the causal factor to the use of AI-enabled legal services. Case narratives for the two typical cases are presented below, including a proposed causal mechanism explaining the link between the identified causal factors and the use of AI.

7.2 Within-case analysis of a typical producer case

Case Narrative of typical producer case (Case 17)

Case 17 describes an individual with several years of experience in a senior producer role. They originally joined Global as a Trainee Solicitor, qualifying into the firm's Banking practice after two years at the firm. As a newly qualified banking lawyer, they gained first-hand experience of AI-enabled legal services in 2017, as a junior team member assigned by their line manager to work on large due diligence exercises for banking clients. Their role within this team was as a consumer of AI-enabled legal services, using Global's own internally developed AI software to augment their work reviewing contracts on behalf of the client. This experience demonstrated to the participant the power of AI to enhance legal work, as previously due diligence tasks were undertaken exclusively by human fee-earners, but the introduction of AI allowed the tasks to be performed quicker, more accurately, and at a lower cost than was previously possible. This led to the participant developing a positive attitude towards the use of AI-enabled legal services.

Following encouragement from senior colleagues, in 2019 the participant became a founder member of a new multi-disciplinary team at Global, which was established to develop technology-driven legal solutions to client problems. This development reflected the firm's belief that technology, and AI in particular, had the potential to significantly improve the delivery of legal services. In this new role the participant moved away from their previous fee-earning responsibilities and began working alongside the firm's technical experts who possessed expert knowledge of AI, but who were not qualified lawyers. Capitalising on their knowledge of banking law and the firm's banking clients, they pooled their legal expertise with the technical expertise of their new colleagues, and transitioned from the role of consumer to producer of AI-enabled legal services.

The producer role the participant now performs required them to develop a much deeper technical understanding of AI software, than is required of someone using the software in a consumer role. Indeed, they are responsible for training consumers how to use the firm's AI software, which means they need to be able to explain and demonstrate how it works to others. They also have project management responsibilities, managing the implementation of technology solutions for clients of the firm. This too requires them to understand how AI software works, and how it is being applied to solve problems on the client's behalf. Finally, they have responsibility for procuring AI software for Global from external technology providers, to complement the AI software that is developed in-house. This requires them to evaluate the suitability of external products to meet the firm's needs. Drawing on these experiences, the process through which the participant became a user of different AI software and integrated them into the firm's own AI-enabled legal services is explained below.

Proposed causal mechanism

The potential causal powers identified in the case relate to both legal professionals as individuals and the AI artifact; meaning a candidate mechanism capable of explaining how some or all of these causal powers were activated, might be found through focusing on the interactions between legal professional and AI. The already completed thematic analysis inferred that questioning practices (Anthony, 2018) could be linked to the development of trust in AI-enabled legal services amongst producers, which might explain AI use amongst

producers. I, therefore, looked for empirical fingerprints that would be suggestive of questioning practices in the data of Case 17. The evidence I identified is detailed below, and suggested to me that questioning is a plausible mechanism that can explain how the case participant determined whether or not to use different AI software. The evidence also suggested that the questioning mechanism operated in distinctive ways, depending on whether the AI software was developed by Global, or procured from an external provider. In describing the proposed causal mechanism I make comparisons with the extant research on questioning where relevant.

Questioning practices take place during human-technology interactions and refer to human behaviours that seek to examine the assumptions that are embedded within technology (Anthony, 2018). The questioning behaviour the participant engaged in when evaluating AI software developed within the firm I call *experimenting*. This behaviour involved the participant's direct involvement in the development and testing of the technology in a 'sandbox' environment using historic client data, in order to determine whether it was suitable to be used in the delivery of actual client work. Experimenting allowed the participant to access the AI's internal model and understand the assumptions on which the model had been built. This made the AI transparent to them and its output explainable.

"There's no AI I'm using right now that is a black box; that I don't know where the answer came from. Even with the machine learning, I know that we have built certain rules, the framework within which it makes decisions. I know what that framework is. And so that if we have seen it has missed something, generally speaking, I can explain why that is."

These experimenting activities appeared similar to those described in the literature as *unfolding* and *framing* (Anthony, 2018). Unfolding is a term used to describe behaviour which allows the hidden features of a technology to be made visible; framing refers to the act of comparing the outputs from a technology to data that has been generated through other sources.

Once AI became transparent, trust was able to develop based on the participant's understanding of the functionality and performance of the technology. This can be seen to reflect the participant's trust in the specific AI software, rather than their trust in the wider

environment in which the technology was used. Experimenting also gave the participant the confidence to defend the performance of the system if it was questioned by others.

“I think [my trust] is based upon the fact that I understand why the system arrived at a particular answer. And I can unpick that, and I can trace that back. So, logic-based systems, or document automation, or the systems where I've given it a set of rules, I can undo these rules, so I always know why it's done something.”

“I'm so familiar with some of the systems I can do it in my sleep...I completely trust it and I know how it works, and I will fully stand behind it if challenged on the output.”

The process through which the participant describes their trust in AI developing reflects their beliefs about the functionality and reliability of the specific AI software, rather than their confidence in institutional sources of trust, such as structural assurances or situational normality.

The participant also recognised experimenting as an approach used by other producers at Global. Experimenting as a questioning practice can, therefore, be understood as a normative behaviour within Global's multi-disciplinary team.

“When I see something work, I trust it...I don't have the classic reservations of, ‘oh, I don't trust it’, or ‘it's a black box’, because my team are a team where we only deploy something if we understand how it works, and we can explain it.”

“I think kind of a bit of a special case is most of my team as we will just give things a bash and we'll just try and work it through. I think a lot of lawyers are just a bit hesitant to do that. They don't understand how to use [AI software] and therefore think, ‘someone else will run it for me’.”

In contrast, experimenting was not understood as characterising the behaviour of legal professionals who were consumers of AI-enabled legal services at Global. In describing the approach of consumers, the case participant stated,

“I think really lawyers [as consumers] will use anything that makes their lives easier, but still enables them to meet their KPIs and meet their targets.”

In describing consumer behaviour in this way, the case participant appeared to be suggesting that consumers would use technology in an uncritical way, provided the output of the technology was understood to help them achieve their professional goals. This description of consumer behaviour is reminiscent of the mechanism of *acceptance*, which describes the output of a technology being treated as legitimate and trustworthy, despite the user lacking understanding of how the technology works or the assumptions embedded within it (Anthony, 2018).

When evaluating technology developed outside the firm by external parties, the participant engaged in a different form of questioning behaviour, which I call *critical examination*. This involved asking the provider of the AI software questions that were designed to generate evidence, which would make the software transparent to the participant. This approach was used when the participant did not have direct access to the inner workings of the technology, thus making experimentation impossible.

“I can be a little bit suspicious when I see third party vendors talking about [their AI system] ‘and it does this, does this, and it does this’. And my questions are, ‘well, that’s great. But how long does it take to train it? How many documents did you need, and how long will it take to deploy it?’”

The participant also recognised the limitations of critical examination as a questioning approach in comparison to experimentation. This reflected their cynicism as to the accuracy and completeness of the information provided by external providers.

“I see technology providers all the time at conferences, pitches, demos and I feel like I go away and never fully understand what they do. I think it’s because I can’t see behind. I can’t see in the black box, I can’t understand why it’s got to a certain point, I need there to be evidence.”

“A big one for me is just sometimes who is explaining it to me, I know that sounds awful, but there are some people I just don’t trust... tech providers need to understand us and what our limitations are, what we’re looking for. And I trust someone who has an honest conversation with me about those limitations as well, rather than promise the whole world.”

To help overcome these concerns, when critical examination was the only means of questioning, the participant indicated they would seek corroborating evidence from other independent sources, in order to further develop their trust in the technology. This might involve asking the AI provider for a client testimonial or looking for evidence that they had previously worked with law firms similar to Global.

“Do you work with other firms like Global? Have you worked on other projects like the projects we work on? What's the feedback been like?...Reviews and previous experience of other people or other businesses is a really helpful benchmark.”

These two acts, which together characterise critical examination, can be understood as complementary to one another, akin to the use of both unfolding and framing, when engaging in experimentation.

In those instances where externally developed software was procured for use by Global, it would subsequently be put through a period of internal testing before being deployed for use by consumers across the firm. This means critical examination was typically used as a preliminary form questioning by producers, prior to them being able to undertake more in-depth experimentation, once the technology had been procured.

In summary, the power of AI to significantly enhance the delivery of legal work, leads to questioning behaviour amongst producers, in order to validate the software's suitability to form part of an AI-enabled legal service. When the software is developed by the firm, questioning takes the form of experimentation (see Figure 14); in contrast, externally developed software is initially questioned through critical-examination, followed by experimentation (see Figure 15).

In both instances successful questioning (part one of the causal mechanism) leads to the validation of the AI software and trust in the AI technology developing (part two of the causal mechanism), the outcome of which is the producer's decision to use AI in the delivery of legal services. This process takes place within a wider market context of law firms seeking to gain a competitive edge over other firms through the use of technology to deliver legal services. Within the law firm itself, the context of the technology pipeline means it is

producers, rather than other legal professionals, who are responsible for validating AI software on the firm’s behalf.

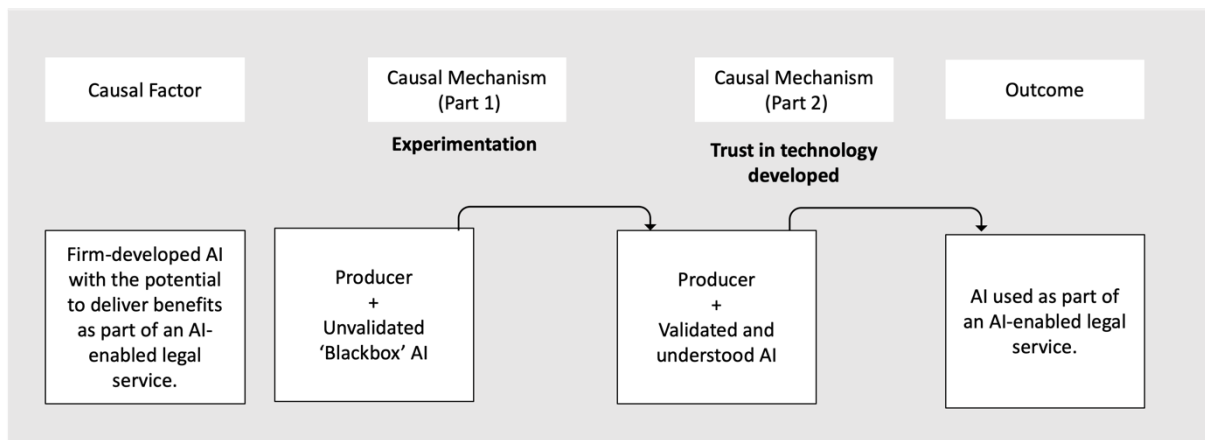


Figure 14: Proposed causal mechanism explaining how producers validate the use of firm-developed AI software.

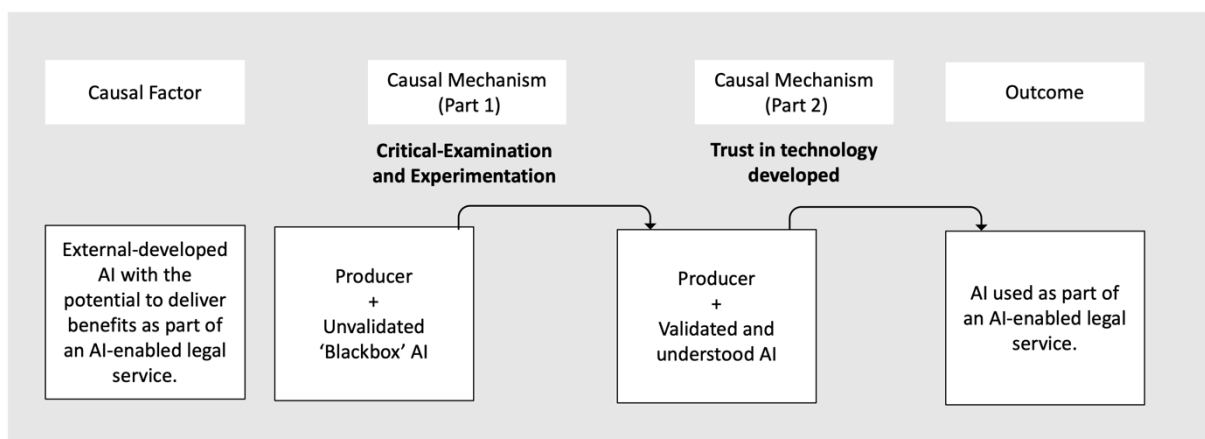


Figure 15: Proposed causal mechanism explaining how producers validate the use of external-developed AI software.

Evaluation of the proposed causal mechanism

In seeking to evaluate the power of the proposed causal mechanism to explain the outcomes of other participants in producer roles, I sought empirical evidence within Cases 1, 6 and 12, which together with Case 17, comprised the original bounded population that was identified. Empirical evidence that was indicative of questioning behaviour (part one of the causal mechanism) and the development of trust in technology (part two of the causal

mechanism), when seen alongside the use of AI within AI-enabled legal services, would be seen to offer cautious support for the proposed causal mechanism. In contrast an absence of evidence for either or both parts of the proposed mechanism, or evidence that could be better explained through reference to an alternative mechanism, would suggest the explanatory power of the proposed causal mechanism did not extend beyond the typical case.

Cases 6 and 12 both contained empirical evidence that was supportive of questioning behaviour involving experimentation rather than critical examination. In Case 6 experimenting could be inferred to take place through the participant's experience of formal training in AI software, and the direct role they played in the development of AI-enabled legal services. These activities required them to look inside the AI software being used at National, which made it understandable to them and explainable to others. From engaging in these behaviours they also indicated that their trust in AI as a technology was enhanced.

"I myself have had a hand in sort of actually building tools or I've seen behind the scenes and also, of course, we're careful with the sort of industry partners that we collaborate with. And we've done work with leading researchers in sort of transparency and explainability. And that gives me a lot of confidence in the tools that we use." [Case 6]

"I think, undertaking a lot of the external training. So I did a lot of stuff. [Named AI provider] have like a great university function. So you can take a few different steps and be guided through how it works. And then I think the feeling on from that, just starting to understand it, realizing that it's perhaps not as complicated as you thought it was, and that actually something like [named AI tool], whilst the output looks very complex, once you start to sort of look behind the scenes and you realise how conditional logic works. And how you can build conditional logic and when you can start to do it yourself and you go, 'OK, I understand this, this is actually very understandable for me'. And whilst at first it was quite an intimidating prospect to understand what it is that this tool does, actually, in reality, I understand it and it kind of enhances my trust in it. I think explainability and understandability is key, isn't it?" [Case 6]

Case 12 explained that they had a healthy scepticism about the performance of technology and that their decision to trust AI was based on their relationship with the software, which reflected their own first-hand experiences. These experiences were described as involving experimenting behaviours, similar to unfolding and framing. In explaining their interactions with AI the participant used the analogy of a car to make their point.

“I think I think it's going to take sort of at least my whole career for there not to be a relationship with [AI] that involves continual curiosity, maintenance, investigation, checking. And, you know, it's kind of like with a car, when we get in the car, we trust that it will get us from A to B, but that's only because we regularly MOT, you know, we have a curious relationship...we have a relationship with its ability to function. We are involved in that. We don't just trust at will, if we haven't MOT'd it for years.” [Case 12]

In explaining how their trust in AI developed, they described trust as being based upon their own direct interactions with the technology. However, their explanation did also recognise that having a relationship with the owner of the technology could also impact their degree of trust, suggesting trust could be developed through relationships with humans that had knowledge of the technology, and not just based on their own first-hand experience of technology.

“I suppose is a qualified yes provided, that I have a detailed relationship with the technology, I think, you know, if I'm thinking about the car, if I have if I'm borrowing someone else's car, I wouldn't necessarily trust it. Unless I knew, unless they have a relationship with me, I wouldn't just jump into a random car.” [Case 12]

The participant in Case 1 was in a more senior role than those in Cases 6 and 12, with responsibility for the procurement of AI software for National. This was reflected in their experiences of engaging with external technology suppliers, through which their questioning behaviour could be inferred. Their description of how they evaluated software suppliers was quite abstract, but both critical examination and experimentation seemed to form part of the process of validation.

“We saw about 10 of the different systems at the same time as [named AI system] who were massively overhyped and couldn't do what they wanted to do. So I think I always end these things with a healthy dose of, you know, scepticism, but a willingness to be proved wrong, and for it to do what it needs to. I'm not coming from a kind of point of I don't believe these tools are ever going to do it...So let's actually see what the functionality is and let's test that ourselves with our use case...I'll remain unconvinced until I've seen kind of what's under the hood.” [Case 1]

The distinction the participant made between the ability of salespeople and technical experts to answer questions about their own AI software does, however, reflect the use of critical examination to gain an understanding of how the technology worked, rather than questions being asked to explore purely commercial considerations, such as price.

“You can soon tell the people who are doing the sales because they're salespeople, or they are actually data scientists very interested in the innovation and the under the hood stuff, because they can answer these kinds of questions and they can foresee them coming kind of stuff, rather than just some shiny sales pitch demo.” [Case 1]

Additionally, it could be inferred from the participant's other responses that they possessed a high-level of technical understanding about how National's AI software worked. This was suggestive of them having undertaken 'experimentation' to learn about the assumptions underpinning the software, although they did not explicitly indicate that this was behaviour, they had personally engaged in.

“So in computer science parlance f score is an average, it's a tried and tested way of measuring the effectiveness of an algorithm matching precision and recall and coming out with the figure. And I suppose because it's on a figure, we know what the maximum is. We know what the minimum is, and we can set where our tolerance for that is, and equally we know that if we gave the same sets of papers to a lawyer, there wouldn't be a hundred percent accuracy.” [Case 1]

In summary, there was some empirical evidence that was supportive of the proposed causal mechanism, across the other three cases within the bounded population. The empirical

evidence within each case was not as rich as the data contained within the typical case (Case 17), from which the mechanism was inferred, meaning it was not always possible to identify each individual element of the proposed causal mechanism, or distinguish its use in validating internal and external AI software. However, as the cases were drawn from both Global and National, there was some evidence to support the existence of the proposed mechanism across different contexts. In conclusion, while the findings should be treated with a degree of caution, they do offer a degree of support for the proposed causal mechanism.

7.3 Within-case analysis of a typical consumer case

Case Narrative of typical consumer case (Case 14)

The participant described in Case 14 shared a number of biographical similarities to Case 17. They too joined Global as a Trainee Solicitor, qualifying in 2017 into the firm's Banking practice after the completion of their two-year training contract, whereupon they began to use AI-enabled legal services to undertake large due diligence exercises. Their role as a junior fee-earner within this team meant they were a consumer of AI-enabled legal services, initially using Global's own AI software to perform contract reviews. The Banking team's use of AI to deliver due diligence work grew during 2018 as the firm's AI software became more advanced, meaning it could process data even more quickly and accurately than before. This allowed the team to conduct their work at a lower cost than competitor firms, who were using more human-intensive methods of data processing. This led to the Banking team winning further client instructions for due diligence and established the team's reputation as the 'go to' firm for AI-enabled legal services amongst banking clients. This high esteem in which the Banking team was held, was something the participant experienced first-hand during their secondment to a leading banking client of the firm, and led to Global being formally recognised as one of the most innovative law firms in Europe. By 2021 the use of AI was established as an intrinsic part of the Banking team's day-to-day client transactions. Together these experiences contributed to the participant developing a positive attitude towards the use of AI-enabled legal services.

The participant is now a senior associate in the Banking team, having worked there for seven years. While their own expertise and responsibilities have grown over time, they remain a consumer of the AI-enabled legal services that Global has implemented to assist their legal professionals in achieving their fee-earning responsibilities. The process through which the participant made AI-enabled legal services part of their professional practice is explained below.

Proposed causal mechanism

In Case 14 there was empirical evidence to suggest that questioning was not a plausible causal mechanism capable of explaining the participant's use of AI-enabled services. In contrast to the producer in Case 17, the participant explained that they had limited understanding of AI and how the software they used actually worked. This was suggestive of the participant not questioning the AI software they use.

"Sorry, I'm not particularly brilliant on the technology speak, so apologies for that... there probably are technical meanings behind [AI] that I don't really know, but I probably see it as a kind of technology solution that's been developed. I know other people use the word robots as AI, erm that's a bit beyond me, I think. But yeah, something some kind of clever technology solution that helps us out, really."

Furthermore, they indicated that they did not believe that it was necessary for them to have a detailed understanding of AI software in their role as a consumer of AI-enabled legal services. This suggested that they did not believe that questioning practices were something they were expected to engage in.

"I don't know if it's [knowledge of AI is] hugely important, erm I'm just a normal lawyer, I'm not in the [Global multi-disciplinary team] or anything like that. And I still think I'm relatively competent and knowledgeable about what they can do for us, without knowing specifically the AI intricacies."

"I can't say I have the opportunity to look under the hood of that really, that's not kind of really in my role. I do understand what it does, kind of, at a high level, and could probably talk at a high level. But in terms of the specifics, probably not."

Finally, they also highlighted what they believed to be normative attitudes of consumers that stopped them engaging in questioning behaviour, which they understood as stemming from the structure of their working environment at Global.

“[There is] an attitude of people being too busy. I think everybody is guilty of that. I definitely am, because it's a big time investment to learn how to use these things properly.

“There is a lot of, you know, fee-earning demands and non fee-earning demands, and adding something that involves maybe watching a training video or doing a little bit of kind of like [your] own personal trial and error work on it. It's a little bit, kind of, another thing to add to the list that people don't necessarily want to do.”

In conclusion, there were a number of pieces of empirical evidence suggesting that the participant did not engage in questioning behaviour, and therefore, this mechanism could not be causally implicated in their use of AI-enabled legal services. I, therefore, sought to identify an alternative causal mechanism that could account for their use of AI.

Trusting and using technology, despite not understanding it, is termed *accepting* behaviour, and has been found to exist amongst skilled professionals (Anthony, 2018; 2021). Given the participant's account indicated that questioning had not taken place, it seemed plausible that accepting was the mechanism that would explain how consumers came to use AI-enabled legal services. I, therefore, sought empirical evidence that would be suggestive of acceptance, with the aim of unpacking the mechanism to understand more precisely how it took place.

In explaining how they came to first use AI-enabled legal services, the participant indicated that the use of AI was encouraged by both senior legal experts within the Banking team and technical experts (producers) at Global.

“AI came to the more junior fee earners, being sold as something that's really quite exciting and great, and it's going to be super useful, from the partner group.”

“All our software solutions, our technology solutions, they come to us through either, you know, our [legal subject matter experts]...or our [multi-disciplinary team]. They often bring these kind of technology ideas and solutions to us.”

This experience suggests that while the use of AI was ultimately a decision to be made by the participant (the use of AI was not made mandatory by the firm) their decision-making process was influenced by the behaviour of senior individuals who held positive views about AI. I, therefore, chose the term *expert endorsement* to describe the mechanism that took place during the interaction between consumers and sponsors of the technology.

This endorsement gave the participant confidence that the AI software had already been tested and approved for use by experts within the firm, meaning they did not feel the need to personally question its use. On the contrary, they understood the use of AI to be in line with the expectations that the firm had for someone in a fee-earning role.

“I’m probably going to be unlikely to be the first person to use it, so I reckon I’d be relatively reliant on somebody else saying, ‘Oh, we have this thing like, it’s good for this.’ Erm yeah, I think that’s something that I can see myself, as not being somebody to be like the first person to use it, if that makes sense, I’d be more reliant on somebody, or someone saying, we have this now, you might like to try this.”

This explanation highlights a temporal dimension to the relationship different people at Global have with AI, with the participant indicating that their first interaction with AI software takes place much later than others, for example those who endorse its use. While not explicitly mentioned in the participant’s account, this temporal ordering could be inferred to reflect the relative position of the different parties to one another, within the context of the technology pipeline of an AI-enabled legal service (Armour and Sako, 2020). With the parties endorsing AI understood to have developed their knowledge and understanding through engaging in questioning behaviour during the initial stages of the pipeline (for example, *Design and procurement*). In contrast, the participant in their role of consumer enters during the final stage (Application of results), which limits their opportunity to engage in questioning behaviour, thus making accepting more likely.

In spite of their limited understanding of the working of AI as a technology, the participant did indicate that they trusted the AI software they used in their work. When explaining how this trust developed, they indicated that this judgement reflected the widespread use of AI software at the firm and the views of those who worked most closely with it.

“Seeing how often we use it, seeing how many things we use it on. Seeing the kind of clients’ success that we seem to have out of it that they can, you know, the main project team that work on it, you know? They all rave about it, and it seems to work, so I guess we’re all kind of more sort of fee earners that help them out. And then we’ve not really got any reason to kind of question its effectiveness...I don’t have any kind of huge, huge trust issues about them really.

The type of trust that emerged from these experiences can be seen to reflect the confidence the participant had in their colleague’s ability to evaluate the effectiveness of AI, in contrast to their, self-acknowledged, limited understanding of AI. Hence, the development of human trust can be seen to follow from expert endorsement and explain the use of AI by those in consumer roles. In addition to this, the participant in referencing the widespread use of AI at the firm, could be understood to be drawing upon an institutional form of trust, reflecting their judgement that the use AI of was in line with their expectations of working at Global i.e. the *situational normality* can be seen to explain the development of trust in AI use (Gefen, Karahanna, and Straub, 2003). On balance, however, the more limited empirical evidence for institutional trust was not in my judgement as persuasive, as the more frequent references to their colleagues as a source of trust.

Interestingly, while human trust in colleagues did lead to the participant’s use of AI-enabled legal services, it may have also impacted their ongoing professional practice. Whereas the use of AI amongst producers reflected their understanding and trust in the technology to generate the correct output; the participant in their role of consumer, indicated an ongoing need to validate the accuracy of the output of the AI-enabled legal services they were using, when the output was not in line with expectations. The precise role they played in validating the output, reflected their level of legal expertise (rather than technical know-how).

In those instances where their legal expertise was limited, they described having responsibility for inputting data into the AI software, but that responsibility for assessing the accuracy of the output was given to a colleague with the relevant legal expertise.

“I've never really kind of have the role of having to interpret the data after that. I'm normally just, you know, the person [who] sort of helps out in the kind of team and fills in the reviews et cetera. I've never, or not for a long time anyway, probably not since 2018, had to kind of be the person that interprets what comes out of it.”

In contrast, in those instances in which, they were tasked with evaluating output accuracy, they compared the actual output of the system to their own expectations, based on their expert wisdom. When discrepancies between the two arose it was assumed an error may have been made, meaning further investigation was required.

“So I know in my mind what this facility agreement should look like. So once I do the questionnaire, I know what should be in the in the output and it's what I'm expecting because it's, you know, its drafting a document that I know relatively inside out.”

In summary, the power of AI to significantly enhance the delivery of legal work, leads to an interaction between consumers of AI and colleagues who the consumer perceives to possess either legal or technical expertise, which makes their assessment of AI influential. When the mechanism of expert endorsement (Figure 16) is successful, consumers treat AI-enabled legal services as having been verified, despite AI remaining a ‘black box’ that is not understood by them. Through developing trust in the human colleagues who endorsed AI, the consumer decides to use AI-enabled legal services as part of their own professional practice. However, this use is characterised by a need for ongoing validation. As before, this process takes place within a wider market context of law firms seeking to gain a competitive edge over other firms through the use of technology to deliver legal services. Within the law firm itself, the technology pipeline model means consumers are not made responsible for validating AI-enabled legal services.

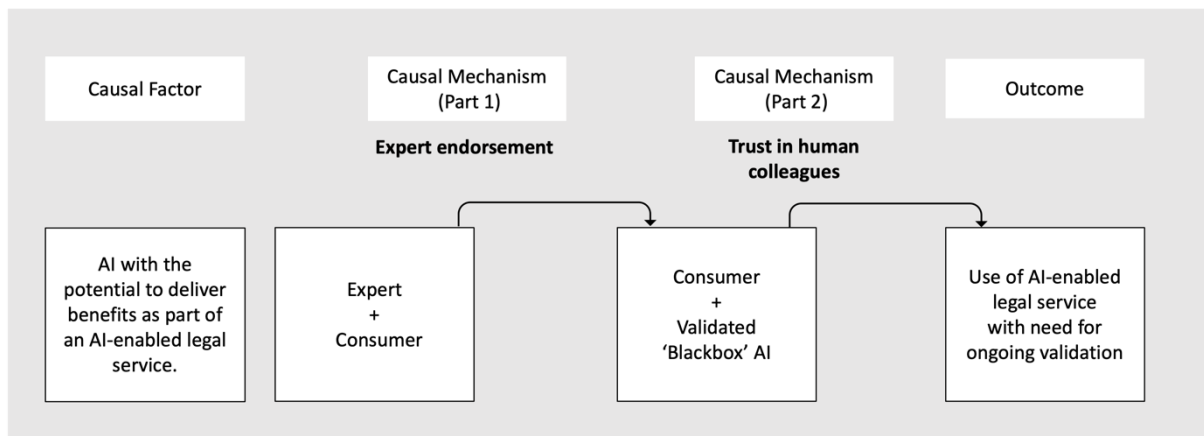


Figure 16: Proposed causal mechanism explaining how consumers decide to use AI-enabled legal services as part of their professional practice.

Evaluation of the proposed causal mechanism

The relevance of the proposed causal mechanism to participants in consumer roles, was evaluated through reference to empirical evidence from the other similar consumer cases that formed the bounded population (Cases 3, 9, 10, 11 and 15). The evaluation, therefore, did not include consumer cases 4, 16 and 18, which fell outside the bounded population because they demonstrated a higher level of similarity to the producer cases.

Empirical evidence that could be traced back to accepting through the mechanism of expert endorsement (part one of the causal mechanism), and the development of trust in human colleagues (part two of the causal mechanism), would offer further support for the explanatory power of the proposed causal mechanism, when found alongside evidence of the use of AI-enabled legal services. Cases in which the above empirical evidence was absent, or alternative mechanisms were seen to offer a more plausible explanation, would be seen to cast doubt on the power of the causal mechanism.

In Case 11 the participant indicated that they would consider AI software endorsed by Global to be more trustworthy. However, this was described at an abstract level as part of a wider discussion about the use of AI software, rather than through reference a specific interaction with colleagues in which expert endorsement took place. The reference to AI that has been used for several years being more trustworthy than new AI, could be understood to reflect the situation normality of the software, an institutional source of

trust. This source of trust was also seen in the typical case, although I judged it to have less explanatory value than colleagues as a source of human trust, per the proposed causal mechanism. Finally, prior to mentioning endorsement, the participant also indicated their desire to understand how AI works; this could be seen to indicate a desire to engage in questioning activity, rather than simply accepting it, casting further doubt on the relevance of the proposed causal mechanism.

“I'd say that's a big part of the trust is understanding the building blocks behind it. Like you say the decision making that's going into it. And then another element of the trust is obviously like, has it been used before? Is it an endorsed tool within the firm or even wider like? There's lots of talk about regulation of AI and that kind of thing, which makes me concerned that, 'oh, maybe it's not regulated'...but if it was, you know, [and] everyone had been using it for years and years, then you'd trust it a lot more than something that's so new.” [Case 11]

In the empirical evidence of Case 15 personal experience of using AI software and its use by competitors, were both highlighted to be important in the development of trust. Rather than evaluating AI's performance based on an understanding of the software itself, the experiences in case 15 highlighted evaluation of the output produced being more important. In addition, it was also recognised that personality traits might affect how an individual interpreted their experiences. No reference was made to either the mechanism of endorsement or human trust in colleagues, meaning neither part of the proposed causal mechanism was supported.

“I think it's I think it's probably predominantly experience. And then part of it is, I think, just your nature your personality, your, you know, your own risk profile if you like...but I think predominantly it's you're experience if you've used this again and again and you see the results, then you build that trust over time.” [Case 15]

“The only other thing is I think that would probably change people's views was if that technology product has been used by other firms that have been successful in using it, that can also add some weight to the argument.” [Case 15]

A third example of personal experience influencing trust development came from Case 10. Here the participant indicated that while they did initially accept the use of AI uncritically when it was presented to them, they also made explicit reference to their experience of validating AI-enabled legal services, through comparing their output against the output of historic processes. This experience can be seen to mirror the ongoing validation identified in the typical case. Interestingly, in this case the empirical evidence suggested that the process of validation undertaken by junior team members was seen to lead to the development of trust amongst their more senior colleagues.

“The proof of the pudding being you need to see some examples of results of it to fully trust it. But I would go in, you know, not necessarily as I full sceptic on it. To use like sort of discussing the rollout of our auto searching tool, we will end up running parallel processes for probably a month to just persuade the partners that it is actually coming up with the same results...but for the most part if someone says, ‘there you go.’ we kind of take it take as read, with the caveat that will always serve lawyers, if something doesn't look quite right, you'll always dig into and try and work out if it is actually correct or not. Yeah, I'd say I'm on sort of upper end of the trusting scale.” [Case 10]

In these three cases the use of AI appeared to be accepted without evidence of either the mechanism of expert endorsement or questioning taking place. No alternative mechanism of acceptance could be readily inferred from the empirical evidence available. Once AI-enabled legal services were being used, there was further evidence that the output of the software, rather than an understanding of its inner workings, was used to maintain trust. This suggests that the (unidentified) causal mechanism that led to the use of AI-enabled legal services was not sufficient to ensure trust was maintained indefinitely.

The empirical evidence in Case 3 referenced the actions of colleagues at National as influential in developing trust in AI, however, the mechanism of expert endorsement was not apparent. Instead, the role of colleagues was better understood as contributing to the development of institutional trust, through activities such as software development and testing, which can be seen to indicate the existence of safeguards (structural assurances) around the use of AI at the firm. Hence, while the use of AI does reflect acceptance rather

than questioning, the mechanism of acceptance cannot be described as expert endorsement, based on the empirical evidence available.

“If we had something [an AI tool] that would sort of suggest legal advice, advice for the client, I think it would depend on sort of who set it up, so if a Partner set it up then I would trust it.” [Case 3]

“So I think and trust because there are people who work in, you know, we have IT, so I trust [AI] because someone else makes sure that it's safe to use, for example, that data is safe...So I trust it because someone else says, you know, there is someone who specializes in it and they set it up.” [Case 3]

The role of institutional safeguards in developing trust in AI was also evident in Case 9, with the participant highlighting that Global's AI-enabled legal services had been designed to minimise the likelihood of human error, something they believed to be more likely than errors arising from AI software. The importance the participant attached to the role of colleagues within this process, is suggestive of the development of human trust, however, this trust reflected their evaluation of their colleague's legal expertise, rather than being linked to the mechanism of expert endorsement. Hence, while this evidence is suggestive of accepting rather than questioning, the proposed causal mechanism is not supported.

“Clients and firms like us very much do rely 100 percent on what output is generated by those kind of programs...the only margin for error is a fault in us in inputting something incorrectly or putting too much emphasis on a certain type of document or certain keyword. And that, again, is something that gets ironed out by virtue of the fact that it's not just me on this case, you know, two or three other lawyers, more senior than me who are also on it, who will go hold on, you haven't thought about this, why haven't you searched for this year or whatever? So as I said, the output from those machines from those programs is 100 percent correct.” [Case 9]

In summary, amongst the other cases within the bounded population there was limited empirical evidence that could be traced to the proposed mechanism of expert endorsement. When combined with the lack of empirical evidence to suggest that

questioning had taken place, I was led to conclude that the use of AI-enabled legal services amongst consumers needed to be explained through an alternative mechanism of acceptance to the one originally proposed in the typical case. Evidence from three of the cases highlighted that the widespread use of AI software made it more trustworthy and more likely to be used. There was also some evidence suggesting that if firms put structural assurances in place around the use of AI, then AI use was likely to take place. Taken together, this case evidence offers cautious support for the existence of a causal mechanism linked to the wider social structure in which AI-enabled legal services exist, rather than a mechanism emerging from the interactions between legal professionals and the AI artifact itself. I also inferred from the data that any alternative mechanism should be capable of explaining the need to perform ongoing validation of AI-enabled legal services, which was apparent in multiple consumer cases.

A further potential explanation consistent with the case evidence is that there is no single causal mechanism capable of explaining the use of AI-enabled legal services amongst consumers, and that there are instead multiple routes through which this outcome can arise. Equifinality is not uncommon in complex social processes, which may also help explain why no bounded populations were identified amongst the environmental causal factors present in Group B. Hence, while the proposed mechanism of expert endorsement was not apparent in the majority of the cases within the bounded population of consumers, its potential causal power should not be discounted in all cases. Indeed, it is worth recalling that in findings of the thematic analysis, which reflected the data of all consumers (rather than the sub-set which formed the bounded population in the process tracing analysis), most consumers indicated that their senior colleagues did influence the development of their initial trust in AI. In conclusion, therefore, expert endorsement may be one of a number of mechanisms capable of explaining the use of AI-enabled legal services amongst consumers, each of which requires a different combination of associated causal factors to be in place, if it is to take effect.

8. DISCUSSION AND RECOMMENDATIONS

The findings of the thematic analysis and process tracing presented in the previous two chapters offer a number of theoretical and practical insights about how legal professionals understand and use AI-enabled legal services as part of their professional practice. In seeking to answer the two research questions that were identified through the literature review, the research makes three distinct contributions, which are relevant to existing conversations about the use of epistemic technologies by knowledge workers and professionals; and the use of AI within the context of a technology pipeline.

The chapter begins by discussing the findings of the research through reference to ongoing theoretical debates relating to knowledge work, AI and the professions. In doing so, I discuss where the findings support, extend and challenge the findings of earlier research, and highlight how my findings can be seen to develop our theoretical understanding of the phenomenon of AI-enabled legal services. Following this, I discuss the implications of the research for organisational practitioners (section 8.4), the limitations of the research (section 8.5) and directions for future research (section 8.6).

8.1 Congruences and incongruences in the understanding of AI-enabled legal services

In addressing the question: *How are AI-enabled legal services understood by the legal professionals that use them?* the findings of this research help distinguish AI-enabled legal services from other technological processes that are used by legal professionals. In addition, the findings also highlight variations in how AI is understood amongst different groups of legal professionals – notably those in producer and consumer roles. This enhances our understanding of AI as a phenomenon, as earlier research on the use of technology in the legal sector has typically focused on describing how technology has impacted the delivery of legal services at both the firm and sector level (Kronblad, 2020; Brooks, Gherhes and Vorley 2020), as opposed to how specific technologies, including AI, are understood by the legal professionals who use them. By shifting the focus to the micro-level to describe how individual legal professionals understand AI-enabled legal services, this research makes it possible to disentangle the use of AI from the use of technology more

generally, while also increasing our awareness of the different factors that can explain AI use amongst legal professionals.

The technological frames that were generated using the data from the thematic analysis reveal the assumptions, expectations and interpretations that legal professionals have about AI-enabled legal services. The structure and content of the frames offer insights into the future behaviour of legal professionals towards AI-enabled legal services, as they indicate which features of the technology and its context legal professionals will focus their attention upon.

Technological frame structure

The findings of this research identify congruence in the overall structure of the technological frames of legal professionals performing a range of roles as National and Global. This suggests that legal professionals develop their understanding of AI through a common set of dimensions. These dimensions were identified as: *personal attitude to AI*, *value of AI*, *firm influence*, *legal sector influence* and *expert influence*. These dimensions highlight how individual understandings of AI are affected by both the technological artifact and the context that surrounds it. In the case of legal professionals, this context includes: individuals with relevant expertise, the law firm they work for, and the wider market for legal services.

The overall structure of the technological frames identified in this research share several similarities with the frames identified in earlier research, which were subsequently used to generate the five-dimension framework proposed by Spieth *et al.* (2021). The findings of my own research, however, differ from Spieth's model in one important respect, which suggests the mental models of AI held by legal professionals are distinctive from those of individuals working in different contexts. Specifically, the dimension Spieth *et al.* (2021) name *supervisor influence* was not found to be salient amongst legal professionals; in its place the novel dimension *expert influence* was evident.

The privileging of expertise reflects an important principle that has been identified as underpinning models of professionalism – that authority is seen to stem from both

expertise and a service ethic. This contrasts with the managerial principles that are more prevalent outside professional service firms, which link authority to accountability and results (Noordegraaf, 2015). For example, in managerial hierarchies supervisors are typically accountable for the actions of their subordinates, and so are given the power to exercise authority over their actions, in contrast professionals often retain a much higher level of personal autonomy. The salience that legal professionals in this research demonstrated towards the views of colleagues with high levels of either technological or legal expertise, rather than their direct supervisors, reflects the ongoing importance of professional logics in explaining the behaviour of legal professionals towards AI.

This structural difference in the technological frame of legal professionals, when compared to workers in other contexts, highlights the importance of adapting existing frameworks to study phenomena within a specific context. Thus, it is suggested that the revised framework proposed in this research, better reflects the factors that explain how legal professionals understand AI-enabled legal services, than the more generic framework proposed by Spieth *et al* (2021). The identified frame structure also acknowledges the distinctive roles of the user, system and organisational context in explaining individual behaviour towards technology, which addresses calls for more inclusive approaches to understanding technology acceptance (Jensen and Aanestad, 2007).

Frame content

The content of the technological frames of the legal professionals in this study also showed evidence of congruence. This indicates that not only do legal professionals pay attention to similar factors when evaluating AI-enabled legal services, but that their understanding of the technology along several of these dimensions is also similar.

The positive *personal attitude to AI* reported by legal professionals in this research reflects their understanding of AI as a tool, the purpose of which is to augment their ability to diagnose legal problems, while leaving them responsible for the 'professional act' of inference making (Abbott, 1988). This interpretation of AI as a largely benign development for legal professionals challenges those narratives that present AI as a threat to the roles and employment of professionals (Frey and Osborne, 2017), and instead supports those

accounts which argue that collaboration between professionals and AI will ultimately drive productivity growth, (Alarie, Niblett and Yoon, 2018; Tredinnick, 2017).

It was also notable that the AI software that featured in this research was considered a 'tool' by legal professionals in a wide range of roles (producers, consumers and liminals). This interpretation indicates that AI is not understood to be categorically different to the other technologies legal professionals use; rather, all tool technologies, including AI, share a common set of assumptions. This can be seen in the expectations held about AI's ability to augment the work of legal professionals, which is reflected in the content of the dimension *value of AI*, where speed and accuracy (rather than insight and knowledge production) are focused on. These expectations are similar to those held about more traditional production technologies. This persistent tool view of AI means, for the moment at least, legal professionals are still some way away from treating AI as an active digital agent (Ågerfalk *et al.*, 2022) and reconceptualising their relationship with AI (see Spring, Faulconbridge and Sarwar, 2022; Gkinko and Elbanna, 2023). The findings, therefore, do not provide supporting evidence to those who suggest significant revisions need to be made to the assumptions that underpin IS theory, when studying AI-related phenomena (Schuetz and Venkatesh, 2020).

There was also evidence of congruence in the content of the frame dimensions that relate to the wider context in which AI is used. The legal professionals in this research indicated that their interpretations of AI-enabled legal services were subject to both *firm influence* and *legal sector influence*. Both these dimensions were seen to positively impact understandings of AI, with both colleagues and clients highlighted as playing distinctive roles. In respect of clients, the findings are characteristic of what Bourmault and Anteby (2023) term 'turning inwards'. This means that while AI is not understood to have changed the fundamental nature of the relationship between legal professionals and their clients, legal professionals do recognise the positive response they receive from clients, when they make AI-enabled legal services part of their professional practice.

Despite the many areas of congruence discussed above, there was evidence of incongruence in the content of the frames of different groups of legal professionals. This can be understood to reflect different experiences of AI amongst the members of each

group. These differences were most evident when comparing the frames of consumers and producers of AI-enabled legal services. This comparison revealed disparities in the content of the dimension *personal attitude to AI*, and the dimension *expert influence*. While ultimately both producers and consumers elected to use AI in their work, the findings of the research indicated that the process through which AI adoption took place differed between the two groups, with producers demonstrating evidence of questioning behaviour, while consumers (and liminals) engaged in acceptance. It is, therefore, suggested that the content of these two dimensions helps explain the observed differences in the behaviour of producers and consumers towards AI-enabled legal services. The key difference being that producers indicate a much greater understanding of AI than consumers. These findings offer support to research within the banking sector conducted by Anthony (2021), and challenge the traditional view that professionals only adopt technology that they have developed a deep understanding of (Bailey and Barley, 2011).

Taken together the above similarities in both the structure and content of the technological frames helps explain why the behaviour of the legal professionals in this study was consistently positive towards AI, with all participants choosing to incorporate the technology into their professional practice. As noted previously, technological frames shape individual behaviour towards technology by focusing the individual's attention on specific factors. The frame structure identified in this research highlights the different factors legal professionals are likely to pay attention to when deciding whether to use AI. The five identified dimensions can be linked to a combination of cognitive processes and social influences, both of which have been demonstrated to affect the perceived usefulness of technology and by corollary its use (Venkatesh & Davis, 2000). Cognitive processes are used to assess the extent to which a technology is capable of helping an individual to perform a task or job role. The dimension *personal attitude to AI* provides an indication of the extent to which a legal professional understands AI to be *relevant* to their work. The dimension *value of AI* reflects a legal professional's perception of how *effectively* AI can perform relevant tasks. In contrast, social influences reflect the external forces that impact an individual's behaviour towards technology. The dimensions *firm influence* and *expert influence* indicate subjective norms about the use AI, i.e. whether there is an expectation amongst other parties that legal professionals should be using AI. Finally, the dimension

legal sector influence reflects norms about AI held by law firm clients, and the extent to which the status of legal professionals is enhanced through their use of AI. In those instances where legal professionals evaluate the content of these dimensions positively, the likelihood that they will use AI increases.

The five-dimension framework identified in this research, therefore, offers a parsimonious model of the factors most likely to affect AI use amongst legal professionals, when compared to context-neutral models of technology acceptance, such as the extended Technology Acceptance Model (TAM) (Venkatesh and Davis, 2000) or the Unified Theory of Acceptance of the use of Technology (UTAUT) (Venkatesh *et al.*, 2003). This judgement reflects the lack of evidence to suggest that all of the constructs within these two models are of relevance to the legal professionals who participated in this research.

More specifically, while the dimension *value of AI* indicates the relevance of TAM's *perceived usefulness* construct (termed *performance expectancy* in UTAUT) to legal professionals' decision to use AI, there was a lack of empirical evidence to suggest that TAM's *perceived ease of use* construct (termed *effort expectancy* in UTAUT) was salient to the decision of legal professionals to use AI. While the absence of empirical evidence relating to *perceived ease of use* does not explain why legal professionals focused on the performance of the system, while appearing insensitive to the degree to which it could make their work more/less difficult, it is plausible that legal professionals have a general expectation that legal work will be complex and challenging, and that the technology they use to undertake such work will display similar characteristics.

The impact that social influences had on legal professionals' use of AI (as captured in the dimensions relating to *firm*, *legal sector* and *experts*) can be understood to reflect TAM's *subjective norm* construct. While this supports the inclusion of this construct in TAM, it also highlights the complexity of the subjective norm construct and the different sources from which norms can arise. This complexity is recognised in UTAUT's construct of *social influence*, which extends TAM's subjective norm construct, to include other related influences such as image and organisational culture. In doing so, UTAUT more accurately captures the different sources of social influence seen in the findings, for example those relating to the dimension *legal sector influence*. However, certain aspects of UTAUT's *social*

influence construct (specifically those relating to the role of supervisors) are not supported by the data generated amongst legal professionals, suggesting that the construct's breadth has led to it containing redundant elements that are not relevant to the local context of legal professionals.

The proposed five-dimension framework generated from the data in this research, therefore, provides a useful starting point for those wishing to conduct similar research, by allowing them to focus on generating and analysing data relating to those factors which are most likely to affect AI use amongst legal professionals, rather than those which have been found to be of limited relevance.

8.2 Extending the concept of consumers and producers of AI-enabled legal services

To date there has been limited opportunity for empirical research to evaluate the conceptual distinction Armour *et al.* (2022) have made between producers and consumers of AI-enabled legal services. By conducting my own research within the context of law firms that utilise technology pipelines to develop AI-enabled legal services, I was able to generate data capable of assessing the utility of the consumer-producer concept in describing the different ways in which legal professionals interact with AI.

Support for the concept of consumers and producers

Armour *et al.* (2022) hypothesise that individuals in consumer roles will require only a limited understanding of how AI works – enough to facilitate their day-to-day use of AI software and explain the output it produces to clients. They, therefore, predict that consumers need only modest increases in their existing knowledge base, and that the value they create as professionals will remain embedded in their legal expertise. The findings of my research support this conceptualisation of the consumer role by demonstrating: the challenges most consumers face when asked to explain their understanding of AI; and the lack of functional transparency consumers attribute to AI (Dispositional Theme 3: Experiential Theme 1.b and 2.b). It was also notable that a number of the consumers at National and Global did not believe that they required a detailed understanding of AI to

perform their role as a legal fee-earner. This suggests a basic level of AI knowledge is sufficient for consumers to interact with AI-enabled legal services at Step 5 of the technology pipeline (application of results), as by this point the service has already been validated by legal professionals in producer roles, operating on the firm's behalf.

While further empirical evidence of Armour *et al.*'s (2022) conceptualisation of the consumer role is required, in order to demonstrate it can be generalised to other legal contexts, the findings of this research offer tentative support for the concept. This is important to our understanding of how legal professionals interact with AI and incorporate it into their professional practice; as the behaviour of consumers challenges earlier theories of technology adoption amongst knowledge workers, which state professionals require a deep understanding of a technology before they are willing to adopt it (Bailey & Barley, 2011). On a practical level, the unanticipated willingness of consumers to accept the use of AI, despite not fully understanding how it works, may help to explain why the legal sector has demonstrated AI adoption levels that are significantly higher than other sectors of the UK economy (Evans and Heimann, 2022), with large commercial firms (some of whom have adopted the technology pipeline model) driving this growth (The Law Society, 2019).

Armour *et al.* (2022) predict that producers, who typically work as part of a multi-disciplinary team, will be drawn from a variety of different professional backgrounds. They indicate that while some producers are likely to have received specialist legal training, others will possess expertise in technical disciplines such as Information Technology and may have no formal legal qualifications. Armour *et al.* (2022) also suggest that producers will require a common vocabulary, and knowledge of one another's disciplines in order to effectively coordinate their work and pool their shared expertise. My empirical findings indicate the existence of multidisciplinary teams (MDTs) at both National and Global comprised of individuals from both legal and technical backgrounds, thus lending support to the conceptualisation of the producer role. These same individuals possessed much more detailed knowledge of AI and displayed confidence in their ability to understand how the AI software at their firm operated (Dispositional Theme 3: Experiential Theme 1.a and 2.a). This finding supports the existence of shared knowledge and a common vocabulary having developed amongst producers, despite their differing training and expertise. This helps

explain why the producers in my research were able to effectively coordinate their work and undertake a variety of tasks across the different steps of the technology pipeline. Indeed, the effectiveness of MDTs in facilitating the introduction of AI-enabled legal services at National and Global, may reflect the existence of professional norms emphasising openness and relationships within the context of MDTs. These norms can be understood as characteristic of connected professionalism (Noordegraaf, 2020); norms which were not evident in the overall participant data (as detailed in Dispositional Theme 2), the majority of whom were consumers with more limited understandings of AI.

A final way in which the empirical data supports the consumer-producer concept, relates to the observed distribution of tasks across the five stages of the technology pipeline. The pipeline model predicts that consumers typically interact with AI in Step 5, whereas producers are involved across all five steps. The role descriptions provided by participants in this research, reflect this understanding of the different ways in which the two groups interact with the technology pipeline.

Extending the concept of consumers and producers: the emergence of the liminal role

While the data generated in my research lend empirical support for the distinction made between consumers and producers, the data also revealed a third group of legal professionals, which I called *Liminals*, operating within the technology pipeline. As this group does not feature in the model proposed by Armour *et al.* (2022), I suggest that the model is extended to more accurately reflect the different roles performed by legal professionals within the pipeline context.

I define liminals as individuals with legal human capital (but not other types of expertise) whose role combines aspects of the work performed by both consumers and producers. This means liminals interact with AI-enabled legal services as both traditional legal fee-earners and as part of a MDT. This distinguishes them from consumers who focus exclusively on fee-earning work, and producers who may not have been legally trained and/or who no longer undertake fee-earning work. Examples of liminals from my research included a qualified solicitor who split their working week between fee-earning (four days) and project work for a MDT (one day); a trainee solicitor who had undertaken two six-

month placements, first as part of a MDT before then moving to a fee-earning role; and a qualified solicitor who was undertaking a full-time, 12-month secondment to a MDT, before returning to fee-earning work. These different examples of the liminal role capture well the 'in-between' nature of the liminal space that these individuals occupy at the jurisdictional boundary that demarcates the fee-earning work of qualified legal professionals from other individuals working within law firms. Linked to this, the MDTs that liminals spent at least part of their time working within, are reflective of the novel structures that have been found to emerge, in order to help individuals navigate liminal spaces (Dale and Burrow, 2007).

The role played by liminals at both National and Global extends across multiple steps of the technology pipeline, but with a principle focus on steps 4 (data labelling) and 5 (application of results), where liminals with legal subject-matter expertise were actively involved in both the training and evaluation of AI models, prior to them being made available to legal fee-earners at the firm. Liminals' involvement in Step 4 of the technology pipeline requires them to be classified as producers using the simpler producer-consumer binary of Armour *et al.* (2022), but unlike producers, liminals did not demonstrate a sophisticated knowledge of AI; they were also likely to indicate that they found AI to lack transparency. This indicates liminals do not possess the shared vocabulary that was evident amongst producers, and may not have yet adopted the same norms of behaviour as other MDT members. These differences have the potential to cause coordination problems when producers and liminals are required to work together on tasks such as data labelling, which may undermine the effectiveness of the technology pipeline and negatively impact levels of trust in AI, due to the role of collaboration within the environmental trust factor (Hancock *et al.*, 2011).

A further liminal responsibility not covered in the original pipeline model was to provide training to consumer colleagues on the use of AI-enabled legal services. Such training can be understood as necessary if fee-earners are to effectively apply the results of an AI-enabled legal service in Step 5 of the pipeline. I hypothesise that liminals may be uniquely well-suited to deliver training to consumers as it will leverage their legal subject-matter expertise, fee-earning experience, and knowledge gained from testing and evaluating the AI software. This may allow them to perform this task more effectively than producers who may lack both formal legal training and experience of fee-earning. *Training delivery* should

therefore be recognised as a specific activity that takes place within the technology pipeline model, between Step 4 (data labelling) and Step 5 (application of results).

In conclusion, while the identification of liminals adds an added layer of complexity to the technology pipeline model, their combination of fee-earning responsibilities with tasks associated with the producer role (in spite of their modest understanding of AI) makes them distinctive from both consumers and producers. What is more, the presence of multiple individuals in liminal roles at both National and Global suggests that the creation of liminal roles is a deliberate choice made by law firms, rather than a localised phenomenon that has arisen by chance. The liminal role is, therefore, worthy of inclusion if the pipeline model is to accurately reflect the empirical reality of organisational life within UK law firms.

8.3 The use of AI-enabled legal services can be explained through two distinct causal pathways

The findings of the process tracing analysis identified two causal pathways through which legal professionals came to use AI-enabled legal services at both National and Global. Through identifying these two pathways my research contributes to the theoretical literature on knowledge work by demonstrating the conditions and mechanisms through which different legal professionals use AI-enabled legal services as part of their professional practice.

In interpreting the data for typical cases of producers and consumers I have theorised how the different roles that legal professionals undertake within the context of a technology pipeline, affects how they make AI-enabled legal services part of their professional practice. More specifically, I have provided insights into how the use of AI by producers can be explained through their adoption of questioning practices, whereas consumers' use of AI is linked to the practice of accepting. These findings add further empirical support to research conducted by Anthony (2018; 2021), which has only recently demonstrated that not all knowledge workers choose to question the epistemic technologies they use as part of their professional practice; and that they are willing to rely on technology to perform professional tasks, without necessarily understanding how the technology works. My research, which

provides empirical evidence of questioning and accepting amongst professionals in the legal sector, therefore, contributes to our theoretical understanding of how knowledge workers respond to new technologies by demonstrating that questioning and accepting practices generalise beyond the finance sector (Anthony, 2021) to professionals working within the legal sector, thus extending the boundary conditions under which the theory has relevance.

The causal pathway of producers

My attempts to unpack the mechanisms of questioning and accepting in the context of the technology pipeline makes visible the specific behaviours that different groups of legal professionals engage in prior to becoming users of AI. Producers were found to engage in behaviours that I describe as *experimentation* and *critical examination*, with the balance between the two reflecting whether they were seeking to validate AI that had been developed internally at their firm or procured from an external party. While experimenting is similar to forms of questioning behaviour reported in other contexts, for example the banking sector (Anthony, 2021); critical examination appears to be a novel behaviour that producers utilise to develop their understanding of technologies that they cannot directly interact with at that time. While sharing some similarities with Anthony's (2021) finding that some senior bankers question AI through critiquing the analysis their more junior colleagues produce using AI (rather than using the technology themselves); critical examination is characterised by producers critiquing the work of experts from other organisations, typically AI software developers. Hence, while Anthony's (2021) research highlights the role that authority derived from an organisation's hierarchy can play in facilitating questioning practices, organisational hierarchy is not implicated in critical examination. From a process perspective, critical examination was also seen to act as a precursor to experimentation, rather than in isolation. This indicates that questioning behaviour can take place over extended periods, during which individuals demonstrate the use of multiple approaches to questioning, which they have developed to meet the needs of different contexts.

The findings of the research also demonstrated that questioning behaviour led to the development of trust in the AI artifact amongst producers. This suggests that the technological dimensions of trust, including the attributes and performance of AI, were

significant to producers' use of AI-enabled legal services. In contrast, there was limited evidence to suggest that the human and environmental factors linked to trust development (outlined by Hancock *et al.*, 2011 and summarised in Figure 7, p.67) were responsible for producers developing trust in AI-enabled legal services. The centrality of the AI artifact in the accounts of how trust develops amongst producers, demonstrates the ongoing relevance of system design features to theories of trust development. Drawing on Li, Hess and Valacich (2008) whose research highlights that different trusting bases will overshadow others depending on the local context, producer trust can be understood to primarily reflect a cognitive calculation that is made about the relevance and quality of the AI artifact, using knowledge which is ascertained through the act of questioning; rather than in light of wider institutional bases relating to the social structure of law firms and the legal profession more generally. This understanding of producer trust as being fundamentally knowledge-based is much simpler than the trust concept detailed in the Trust-TAM model (Gefen, Karahanna, and Straub, 2003), which identifies multiple trust antecedents. While knowledge-based trust is one of these, the antecedents of *structural assurances* and *situational normality* (both institutional bases) were not found to be relevant to trust development amongst producers. Hence, in seeking to understand the role of trust in producers' use of AI, trust should be narrowly conceptualised as knowledge-based trust.

The causal pathway of consumers

While evidence of both experimentation and critical examination was found in the majority of producer cases, there was a lack of evidence to suggest consumers engaged in similar practices. This provides support to the claim that some legal professionals will use AI, despite not understanding how it works, which is characteristic of individuals who engage in accepting behaviour (Anthony, 2018).

The analysis of accepting practices amongst consumers did not identify any single mechanism across the majority of the consumer cases analysed. While there remains the possibility of an as yet unidentified causal mechanism that is common to the majority of consumers, I judged it more likely that multiple mechanisms of acceptance exist alongside one another. Such equifinality is not uncommon in complex socio-technical systems

(Henfridsson and Bygstad, 2013), and may reflect the greater diversity of job roles and experiences amongst consumers, when compared to producers.

It was also possible to identify from the data that consumers' trust in AI was not based on the technology-related factors that explained the development of knowledge-based trust in producers. Hence, while consumers did indicate that their decision to use AI-enabled legal services required them to develop trust in the service, the process through which trust developed was distinctive. Amongst the analysed consumer cases, environmental trust factors, including institutional sources of trust were in evidence, indicating an important difference in the causal pathways of producers and consumers. Consumers can, therefore, be understood to draw on a different set of trusting bases, despite working in a similar organisational context to producers. This is reflected in consumer accounts which highlighted their trust in those colleagues who possessed significant legal or technical expertise; and their knowledge of the institutional safeguards put around AI by their firm. Consumer trust can, therefore, be understood to combine the two institutional antecedents of the Trust-TAM model (Gefen, Karahanna, and Straub, 2003) – *structural assurances* and *situational normality* – neither of which were implicated in trust development amongst producers.

In conclusion, by looking at the two causal pathways together (rather than in isolation from one another) the retention of the multi-antecedent conceptualisation of trust used in Trust-TAM is recommended, in order to fully capture the distinctive ways in which trust in AI develops amongst consumers and producers of AI-enabled legal services.

Linking the causal pathways to the technology pipeline

The different steps of the technology pipeline at which consumers and producers come into contact with AI offers a further opportunity to theorise why the above two causal pathways have developed. The extensive involvement of producers across all five stages of the technology pipeline gives them the time and exposure to AI that is required to engage in questioning practices. This allows them to develop their knowledge and understanding of AI, which subsequently makes it possible for them to develop knowledge-based trust through their own evaluation of AI as a technological artifact. Temporally the position of

producers within the pipeline also means that they are likely to be among the first people within a law firm who come into contact with new AI software, meaning the process of trust development can start earlier than it would for a consumer.

In contrast, consumers are restricted to the final stage of the technology pipeline, which when combined with the heavy demands of their fee-earning work, means they have less time to interact with AI software, and when they do, their interactions are more restricted. By not being exposed to the design and testing phases of the pipeline (Steps 2 to 4), consumers miss out on the opportunity to develop a detailed knowledge of the inner workings of AI, meaning AI remains 'black boxed'. This explains why the initial trust of consumers does not develop through the knowledge-based trust in technology pathway seen with producers; instead trust is likely to develop indirectly, either through interactions with colleagues at Stage 5 of the pipeline, or through an awareness of the institutional safeguards surrounding the use of AI, which have been established by producers during earlier stages of the pipeline.

The impact that the technology pipeline has on task allocation amongst legal professionals, and through this opportunities to engage in questioning and accepting behaviours, helps to explain why certain 'targets of trust' become most salient to different individuals (Söllner, Hoffmann and Leimeister, 2016). The effect of which is ultimately seen through the different ways in which professionals respond to the introduction of new workplace technologies, such as AI.

8.4 Implications for organisational practitioners

The following recommendations are directed to organisational practitioners working in large UK law firms whose strategy involves a significant commitment to the use of AI-enabled legal services. The recommendations are designed to ensure the use of AI is for the benefit of both the firm and the legal professionals working within it. The recommendations are not to be generalised beyond this context to organisations in other sectors and should not

be regarded as an endorsement of the use of AI-enabled legal services, as such a judgement can only be made through reference to the specific goals and context of the firm.

- 1. Firms seeking to promote the use of AI-enabled legal services should adopt a multi-dimensional approach.** The positive attitudes to AI-enabled legal services found amongst the legal professionals in this research reflected not just their understanding of the technology and its capabilities, but also wider social influences, including the opinions of colleagues and clients of the firm. As the decision to use AI-enabled legal services is likely to be shaped by all five dimensions of the technological frame identified in this research, firms wanting to maximise AI use should look to address each of these when implementing AI-enabled legal services. A successful programme of implementation should therefore address: the relevance of AI to the work of legal professionals; the specific ways in which AI can create value; wider stakeholder views of AI, in particular those held by clients and colleagues who are widely recognised as legal or technical experts.
- 2. Multi-disciplinary teams can support the introduction of AI-enabled legal services.** The findings of this research and other recent empirical studies (Armour, Parnham and Sako, 2022) demonstrate firms with multi-disciplinary teams (MDTs) have been able to successfully implement AI enabled legal services. While a causal link between MDTs and the success of AI-enabled services still needs to be demonstrated, it seems probable that such teams offer an effective way to connect professionals from different disciplinary backgrounds, so that they can collaborate in the creation and application of AI-enabled legal services. The creation of MDTs is, therefore, worthy of consideration when firms are planning to introduce AI-enabled legal services.
- 3. Firms need to think carefully before creating liminal roles.** The research showed that legal professionals in liminal roles have an understanding of AI that is similar to consumers, despite having responsibility for work associated with the producer role. Individuals moving from consumer roles into liminal roles, therefore, need clarity about the role they are expected to perform, and are likely to require an enhanced level of training designed to further develop their understanding of AI, so it more closely

resembles that of producers. A formal selection process for liminal roles would allow an applicant's suitability for a liminal role to be evaluated consistently. The use of a training needs analysis would pinpoint individuals who would benefit from: developing their general knowledge of AI; learning the methodologies used by MDTs (such as the 'technology pipeline'); understanding the precise nature of the role played by liminals at the firm. This training should help liminal members of a MDT to converse with other team members using a common vocabulary, thereby allowing them to take part in a wider range of activities e.g. *requirements gathering*, rather than being limited to activities that only require legal expertise e.g. *data labelling*.

4. **'AI champions' could offer a credible source of human trust amongst consumers of AI-enabled legal services.** The findings of this research indicate that some consumers develop their initial trust in AI based on their existing network of contacts. In order to ensure that all legal professionals have credible sources of AI knowledge available to them, 'AI Champions' should be identified and made visible to potential consumers of AI-enabled legal services. This could be achieved by asking each of the firm's practice areas to identify a small number of individuals who are widely respected for their legal expertise and fee-earning ability, and who are also experienced users of AI-enabled legal services. This combination of qualities is likely to make their views of AI more influential amongst their close colleagues and minimise the likelihood that legal professionals conflate seniority with AI subject-matter expertise. AI champions should be made available to answer questions and address any concerns amongst members of their practice area who are not yet users of AI-enabled legal services. This could be achieved using a combination of pre-recorded digital content to be distributed on the firm's intranet and synchronous questions and answers sessions, designed to meet the needs of hybrid professionals.
5. **AI transparency supports the development of trust in AI technology.** Legal professionals who can understand how AI technology works are able to perform a wider range of tasks within a technology pipeline and can develop trust in AI independently of their colleagues. A firm-wide commitment to AI transparency would allow all legal professionals to better understand AI, without having to generate this knowledge

themselves through questioning practices. Understanding how AI works would empower them to make informed decisions about their use of AI, rather than making them reliant upon the knowledge of others, or firm-wide safeguards for AI use. AI transparency would also promote the development of a firm-wide technological frame and reduce the likelihood of content incongruency developing between different groups of legal professionals, thus making behaviour towards AI more consistent across the firm. It would also enhance legal professionals' ability to explain the outputs of AI-enabled legal services to their clients. AI transparency could be promoted by providing legal professionals – particularly those in consumer and liminal roles – with a non-technical explanation of the overall purpose of the service and how the different AI elements of the process work. Explanations should be short and use language that can be easily understood by legal professionals and clients of the firm. Legal professionals in producer roles would be well-placed to take responsibility for promoting AI transparency across the firm.

8.5 Limitations of the research

The research design I chose for this research placed limits on the types of claims that it would be possible for me to make based on my analysis of the data. The limitations identified below are not, however, a repeat of the limitations that were discussed in the methodology chapter, and accepted as inherent to the way in which the research was conducted e.g. limited scope to generalise the results beyond the context of UK law firms. Rather the issues identified below relate to the developments in AI technology that have taken place during the time it has taken to complete my research, or are methodological limitations that could potentially have been overcome with a different research design. The findings of my research, therefore, require critical interpretation in light of the following limitations:

1. During the course of my research AI technology has developed rapidly. A significant change is the now widespread availability of generative AI models, such as Chat GPT, which are now beginning to be used by a number of organisations, including law firms. Generative AI has the potential to impact the work of professionals in ways that the AI

software that featured in this research was unable to, for example a greater role in tasks that require inference making. It is, therefore, possible that the assumptions and expectations that legal professionals have about AI, will have shifted in fundamental ways since the research was completed.

2. By allowing participants to be identified through self-selection the participant sample cannot be understood to reflect the wider population of individuals at Global and National who have experience of using AI-enabled legal services. This means the participant responses on which my analysis was based, are highly unlikely to reflect the full range of views held about AI-enabled legal services, meaning the results present a partial understanding of the phenomenon. The most striking example of this is the lack of any participants who had a negative attitude towards AI-enabled legal services, or who had discontinued using them. A direct impact of this was my inability to make my anticipated comparison of users and non-users of AI-enabled legal services. With the benefit of experience I think this limitation was to an extent avoidable had I been clearer in my approach to communicating who was eligible to take part in the research.
3. The contribution made by qualitative research is heavily dependent upon the quality of the inferences the researcher is able to draw during their analysis. While I was alert to this prior to starting my own analysis, I became very aware of my limited interpretive repertoire (Alvesson and Sköldbberg, 2018) when generating dispositional themes for the thematic analysis and identifying causal mechanisms during process tracing. While this limitation is one that is likely to be true of most junior researchers, my decision to employ a retroductive approach to theorising compounded the issue; as rather than encouraging me to analyse my data using a pre-selected theoretical framework (that I could have got to know intimately), the approach rewards researchers who are able to draw upon ideas from a wide range of theoretical perspectives. With the benefit of experience, I still see the merits in such an approach, as it offers the potential for a more comprehensive understanding of the phenomenon, however, knowledge of alternative theoretical perspectives may have generated further insights.

8.6 Directions for future research

There are several interesting opportunities for future research about the use of AI by professionals.

1. My analysis suggested that legal professionals who use AI-enabled legal services within the context of large UK commercial law firms share much in common, in terms of the mental models (technological frames) they use to understand AI, but that important differences can arise between legal professionals in different roles. Future research could investigate the technological frames of professionals who use AI in law firms that have not adopted the technology pipeline model, or who work for firms in different legal jurisdictions, to see whether this changed context leads to the development of a distinct technological frame. The content and structure of the frames of legal professionals who do not use AI as part of their professional practice also merit research. This might shed light on any differences in the assumptions and expectations that such professionals hold about the technology and help to explain their decision not to use AI-enabled legal services. Ultimately, as research in this area matures the opportunity will arise to synthesise the available empirical data using techniques similar to the factor analysis performed by Spieth *et al.* (2021). This would allow the development of a measurement instrument capable of generating quantitative data about the technological frames of legal professionals working across different legal settings.
2. This research offers the category of liminal as an addition to the previously identified categories of producer and consumer. This proposed extension to Armour *et al.*'s (2022) conceptualisation of the technology pipeline should now be investigated further to determine the extent to which it is a local phenomenon or reflective of other contexts in which the pipeline approach has been implemented. It is recommended that future research within this context considers the existence of all three roles when analysing the membership of multi-disciplinary teams. This research could initially be undertaken in the context of law firms, before looking for evidence of liminals in other industry sectors in which AI is being implemented, for example accountancy and management consultancy. In the event that liminals are found to be a typical feature of MDTs, future research could look to further distinguish the roles of consumers, producers and

liminals. Research could also be undertaken to investigate the reasons for liminal roles existing. For example, does the liminal role play an important role in coordinating the work of producers and consumers? The use of longitudinal research methods would also allow the experiences of individuals in liminal roles to be tracked over time. This would have the potential to reveal the existence of 'rites of passage' (Turner & Abrahams, 2017) associated with liminal roles. For example, do individuals in liminal roles typically transition to producer roles, return to consumers roles, or retain their liminal status?

3. A surprising finding of this research was that all of the legal professionals who volunteered to take part were users of AI-enabled legal services. As acknowledged above this finding likely reflects methodological weaknesses in this research; indeed it was referenced by participants that not all legal professionals in their firm choose to use AI-enabled legal services when offered them. Future research should, therefore, focus on generating data that allow for a comparison to be made between the users and non-users of AI-enabled legal services, in order to better understand the reasons for use and non-use of AI-enabled legal services. The absence of non-users of AI also raises the question of what explains the apparent reticence of these individuals to share their views about this topic. The strategic use of rhetoric has been identified as a strategy that can be used to shift institutional logics and secure new settlements within organisations (Suddaby and Greenwood, 2005). There is, therefore, an opportunity for future research to investigate whether management discourses extolling the use of AI within law firms has effectively silenced individuals who hold alternative views about this phenomenon.
4. The findings of this research were hampered by a lack of more detailed and voluminous empirical data. There is, therefore, an opportunity to undertake future research using methodologies that were not available to researchers during the period affected by the COVID pandemic. An ethnographic approach would allow data to be generated using a wider range of methods e.g. regular observations of legal professionals as they use AI-enabled services. Ethnographic approaches would also enable longitudinal data to be collected, for example data could be generated with participants at different time points

during the period in which AI-enabled legal services are implemented. This type of data would enrich the insights that can be offered by analytical approaches such as process tracing, as the temporal aspect of the data would make the tracing process easier. Longitudinal research would also allow changes in the structure and content of a legal professional's technological frames to be examined.

REFERENCES

- Abbott, A. (1988) *The System of Professions: An Essay on the Division of Expert Labor*. Chicago, IL: University of Chicago Press. Available at: <https://press.uchicago.edu/ucp/books/book/chicago/S/bo5965590.html> (Accessed: 21 September 2023).
- Acemoglu, D. and Restrepo, P. (2018) 'The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment', *American Economic Review*, 108(6), pp. 1488–1542. Available at: <https://doi.org/10.1257/aer.20160696>.
- Ackroyd, S. and Muzio, D. (2007) 'The Reconstructed Professional Firm: Explaining Change in English Legal Practices', *Organization Studies*, 28(5), pp. 729–747. Available at: <https://doi.org/10.1177/0170840607073077>.
- Adams, C. and Thompson, T.L. (2016) *Researching a Posthuman World: Interviews with Digital Objects*. London: Palgrave Macmillan UK. Available at: <https://doi.org/10.1057/978-1-137-57162-5>.
- Adams, T.L. (2020) 'Professional employees and professional managers: conflicting logics, hybridity, and restratification', *Journal of Professions and Organization*, 7(1), pp. 101–115. Available at: <https://doi.org/10.1093/jpo/joaa005>.
- Ågerfalk, P.J. *et al.* (2022) 'Artificial Intelligence in Information Systems: State of the Art and Research Roadmap', *Communications of the Association for Information Systems*, 50(1), pp. 420–438. Available at: <https://doi.org/10.17705/1CAIS.05017>.
- Ajzen, I. (1985) 'From Intentions to Actions: A Theory of Planned Behavior', in J. Kuhl and J. Beckmann (eds) *Action Control: From Cognition to Behavior*. Berlin, Heidelberg: Springer (SSSP Springer Series in Social Psychology), pp. 11–39. Available at: https://doi.org/10.1007/978-3-642-69746-3_2.
- Alarie, B., Niblett, A. and Yoon, A.H. (2018) 'How artificial intelligence will affect the practice of law', *University of Toronto Law Journal*, 68(supplement 1), pp. 106–124. Available at: <https://doi.org/10.3138/utlj.2017-0052>.
- Aldammagh, Z., Abdeljawad, R. and Obaid, T. (2021) 'Predicting Mobile Banking Adoption: An Integration of TAM and TPB with Trust and Perceived Risk', *Financial Internet Quarterly*, 17(3), pp. 35–46. Available at: <https://doi.org/10.2478/fiqf-2021-0017>.
- Allen, D.K. *et al.* (2013) 'How Should Technology-Mediated Organizational Change Be Explained? A Comparison of the Contributions of Critical Realism and Activity Theory', *MIS Quarterly*, 37(3), pp. 835–854. Available at: <https://doi.org/10.25300/MISQ/2013/37.3.08>.
- Alvesson, M. (2001) 'Knowledge Work: Ambiguity, Image and Identity', *Human Relations*, 54(7), pp. 863–886. Available at: <https://doi.org/10.1177/0018726701547004>.
- Alvesson, M. (2011) *Interpreting Interviews*. SAGE Publications Ltd. Available at: <https://doi.org/10.4135/9781446268353>.
- Alvesson, M. and Sköldbberg, K. (2018) *Reflexive methodology: new vistas for qualitative research*. Third edition. Los Angeles: Sage.

- Alvesson, M. and Willmott, H. (2002) 'Identity Regulation as Organizational Control: Producing the Appropriate Individual', *Journal of Management Studies*, 39(5), pp. 619–644. Available at: <https://doi.org/10.1111/1467-6486.00305>.
- Anteby, M., Chan, C.K. and DiBenigno, J. (2016) 'Three Lenses on Occupations and Professions in Organizations: *Becoming, Doing, and Relating*', *Academy of Management Annals*, 10(1), pp. 183–244. Available at: <https://doi.org/10.5465/19416520.2016.1120962>.
- Anthony, C. (2018) 'To Question or Accept? How Status Differences Influence Responses to New Epistemic Technologies in Knowledge Work', *Academy of Management Review*, 43(4), pp. 661–679. Available at: <https://doi.org/10.5465/amr.2016.0334>.
- Anthony, C. (2021) 'When Knowledge Work and Analytical Technologies Collide: The Practices and Consequences of Black Boxing Algorithmic Technologies', *Administrative Science Quarterly*, 66(4), pp. 1173–1212. Available at: <https://doi.org/10.1177/00018392211016755>.
- Anthony, C., Bechky, B.A. and Fayard, A.-L. (2023) "'Collaborating" with AI: Taking a System View to Explore the Future of Work', *Organization Science*, 34(5), pp. 1672–1694. Available at: <https://doi.org/10.1287/orsc.2022.1651>.
- Archer, M.S. (1995) *Realist social theory: the morphogenetic approach*. Cambridge, United Kingdom: Cambridge University Press.
- Archibald, M.M. *et al.* (2019) 'Using Zoom Videoconferencing for Qualitative Data Collection: Perceptions and Experiences of Researchers and Participants', *International Journal of Qualitative Methods*, 18, p. 1609406919874596. Available at: <https://doi.org/10.1177/1609406919874596>.
- Armitage, C.J. and Conner, M. (2001) 'Efficacy of the Theory of Planned Behaviour: A meta-analytic review', *British Journal of Social Psychology*, 40(4), pp. 471–499. Available at: <https://doi.org/10.1348/014466601164939>.
- Armour, J., Parnham, R. and Sako, M. (2022) 'Augmented Lawyering', *University of Illinois Law Review*, 2022(1), pp. pp71-138.
- Armour, J. and Sako, M. (2020) 'AI-enabled business models in legal services: from traditional law firms to next-generation law companies?', *Journal of Professions and Organization*, 7(1), pp. 27–46. Available at: <https://doi.org/10.1093/jpo/joaa001>.
- Arntz, M., Gregory, T. and Zierahn, U. (2016) *The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis*. Paris: OECD. Available at: <https://doi.org/10.1787/5jlz9h56dvq7-en>.
- Atkinson, P. and Silverman, D. (1997) 'Kundera's Immortality: The Interview Society and the Invention of the Self', *Qualitative Inquiry*, 3(3), pp. 304–325. Available at: <https://doi.org/10.1177/107780049700300304>.
- Aulakh, S. *et al.* (2017) *Mapping advantages and disadvantages: Diversity in the legal profession in England and Wales*. Available at: <https://www.sra.org.uk/sra/research-publications/diversity-legal-profession/> (Accessed: 25 September 2023).
- Avenier, M. and Thomas, C. (2015) 'Finding one's way around various methodological guidelines for doing rigorous case studies: A comparison of four epistemological frameworks', *Systèmes d'information & management*, 20(1), p. 61. Available at: <https://doi.org/10.3917/sim.151.0061>.

- Bailey, D.E. and Barley, S.R. (2011) 'Teaching-Learning Ecologies: Mapping the Environment to Structure Through Action', *Organization Science*, 22(1), pp. 262–285. Available at: <https://doi.org/10.1287/orsc.1090.0511>.
- Bailey, D.E. and Barley, S.R. (2020) 'Beyond design and use: How scholars should study intelligent technologies', *Information and Organization*, 30(2), pp. 1–12. Available at: <https://doi.org/10.1016/j.infoandorg.2019.100286>.
- Bailey, D.E., Leonardi, P.M. and Barley, S.R. (2012) 'The Lure of the Virtual', *Organization Science*, 23(5), pp. 1485–1504. Available at: <https://doi.org/10.1287/orsc.1110.0703>.
- Baker, S.E. and Edwards, R. (eds) (2012) 'How many qualitative interviews is enough? Expert voices and early career reflections on sampling and cases in qualitative research'. National Centre for Research Methods Reviews.
- Bamberger, P.A. (2018) 'AMD—Clarifying What We Are about and Where We Are Going', *Academy of Management Discoveries*, 4(1), pp. 1–10. Available at: <https://doi.org/10.5465/amd.2018.0003>.
- Bandura, A. (1977) 'Self-efficacy: Toward a unifying theory of behavioral change', *Psychological Review*, 84(2), pp. 191–215.
- Barrett, M. *et al.* (2012) 'Reconfiguring Boundary Relations: Robotic Innovations in Pharmacy Work', *Organization Science*, 23(5), pp. 1448–1466. Available at: <https://doi.org/10.1287/orsc.1100.0639>.
- Beach, D. and Pedersen, R.B. (2016) *Causal Case Study Methods: Foundations and Guidelines for Comparing, Matching, and Tracing*. University of Michigan Press.
- Bechky, B.A. (2006) 'Talking About Machines, Thick Description, and Knowledge Work', *Organization Studies*, 27(12), pp. 1757–1768. Available at: <https://doi.org/10.1177/0170840606071894>.
- Bell, F., Rogers, J. and Legg, M. (2019) 'Lawyer Wellbeing in the (Robotic) Face of Technological Change'. Rochester, NY. Available at: <https://papers.ssrn.com/abstract=3589151> (Accessed: 23 September 2023).
- Berente, N. *et al.* (2019) 'Call for papers MISQ special issue on managing AI.', *MIS Quarterly*, p. pp.1-5.
- Bévort, F. and Suddaby, R. (2016) 'Scripting professional identities: how individuals make sense of contradictory institutional logics', *Journal of Professions and Organization*, 3(1), pp. 17–38. Available at: <https://doi.org/10.1093/jpo/jov007>.
- Bhaskar, R. (1997) *A realist theory of science*. London, United Kingdom: Verso (Verso classics, 9).
- Bhaskar, R. (1998) *The possibility of naturalism: a philosophical critique of the contemporary human sciences*. 3rd ed. London ; New York: Routledge.
- Blumer, H. (1954) 'What is Wrong with Social Theory?', *American sociological review*, 19(1), pp. pp3-10. Available at: <https://doi.org/10.2307/2088165>.
- Boell, S. and Cecez-Kecmanovic, D. (2014) 'A Hermeneutic Approach for Conducting Literature Reviews and Literature Searches', *Communications of the Association for Information Systems*, 34(1). Available at: <https://doi.org/10.17705/1CAIS.03412>.

- Bogner, A., Littig, B. and Menz, W. (eds) (2009) *Interviewing Experts*. London: Palgrave Macmillan UK. Available at: <https://doi.org/10.1057/9780230244276>.
- Bourmault, N. and Anteby, M. (2023) 'Rebooting One's Professional Work: The Case of French Anesthesiologists Using Hypnosis', *Administrative Science Quarterly*, 68(4), pp. 913–955. Available at: <https://doi.org/10.1177/00018392231190300>.
- Bowen, G.A. (2006) 'Grounded Theory and Sensitizing Concepts', *International Journal of Qualitative Methods*, 5(3), pp. 12–23. Available at: <https://doi.org/10.1177/160940690600500304>.
- Bowles, J. (2014) 'The computerisation of European jobs', *Bruegel | The Brussels-based economic think tank*, 24 July. Available at: <https://www.bruegel.org/blog-post/computerisation-european-jobs> (Accessed: 21 September 2023).
- Boyatzis, R.E. (1998) *Transforming qualitative information: Thematic analysis and code development*. Thousand Oaks, CA, US: Sage Publications, Inc (Transforming qualitative information: Thematic analysis and code development), pp. xvi, 184.
- Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3(2), pp. 77–101. Available at: <https://doi.org/10.1191/1478088706qp063oa>.
- British Psychological Society (2021a) *BPS Code of Human Research Ethics*, BPS. Available at: <https://www.bps.org.uk/guideline/bps-code-human-research-ethics> (Accessed: 23 September 2023).
- British Psychological Society (2021b) *Ethics guidelines for internet-mediated research*, BPS. Available at: <https://www.bps.org.uk/guideline/ethics-guidelines-internet-mediated-research> (Accessed: 23 September 2023).
- Brock, D.M., Leblebici, H. and Muzio, D. (2014) 'Understanding professionals and their workplaces: The mission of the Journal of Professions and Organization', *Journal of Professions and Organization*, 1(1), pp. 1–15. Available at: <https://doi.org/10.1093/jpo/jot006>.
- Brönnimann, A. (2022) 'How to phrase critical realist interview questions in applied social science research', *Journal of Critical Realism*, 21(1), pp. 1–24. Available at: <https://doi.org/10.1080/14767430.2021.1966719>.
- Brooks, C., Gherhes, C. and Vorley, T. (2020) 'Artificial intelligence in the legal sector: pressures and challenges of transformation', *Cambridge Journal of Regions, Economy and Society*, 13(1), pp. 135–152. Available at: <https://doi.org/10.1093/cjres/rsz026>.
- Brynjolfsson, E. and McAfee, A. (2014) *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. New York London: W.W. Norton & Company.
- Bucholtz, M. (2000) 'The politics of transcription', *Journal of Pragmatics*, 32(10), pp. 1439–1465. Available at: [https://doi.org/10.1016/S0378-2166\(99\)00094-6](https://doi.org/10.1016/S0378-2166(99)00094-6).
- Butcher, Fiona D. 2022. "Psycho-social Factors Influencing Trust in Artificial Intelligence Advice Systems." Thesis. Leicester: University of Leicester. doi:10.25392/leicester.data.20310096.v1
- Callegari, B. and Rai, R.S. (2022) 'Digitalization and law: innovating around the boundaries', *Journal of Professions and Organization*, 9(3), pp. 291–302. Available at: <https://doi.org/10.1093/jpo/joac019>.

- Cave, S. *et al.* (2018) 'Portrayals and perceptions of AI and why they matter'. Available at: <https://www.repository.cam.ac.uk/handle/1810/287193> (Accessed: 21 September 2023).
- Çelebi, H.İ. (2021) 'Artificial Intelligence Applications in Management Information Systems: A Comprehensive Systematic Review with Business Analytics Perspective', *Artificial Intelligence Theory and Applications*, 1(1), pp. 25–56.
- Charmaz, K. (2006) *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*. SAGE.
- ChatGPT (no date). Available at: <https://chat.openai.com> (Accessed: 21 September 2023).
- Chatterjee, S. *et al.* (2021) 'A possible conceptualization of the information systems (IS) artifact: A general systems theory perspective¹', *Information Systems Journal*, 31(4), pp. 550–578. Available at: <https://doi.org/10.1111/isj.12320>.
- Chen, Y. and Reay, T. (2021) 'Responding to imposed job redesign: The evolving dynamics of work and identity in restructuring professional identity', *Human Relations*, 74(10), pp. 1541–1571. Available at: <https://doi.org/10.1177/0018726720906437>.
- Chreim, S., Williams, B.E. (Bernie) and Hinings, C.R. (Bob) (2007) 'Interlevel Influences on the Reconstruction of Professional Role Identity', *Academy of Management Journal*, 50(6), pp. 1515–1539. Available at: <https://doi.org/10.5465/amj.2007.28226248>.
- Clifton, J., Glasmeier, A. and Gray, M. (2020) 'When machines think for us: the consequences for work and place', *Cambridge Journal of Regions, Economy and Society*, 13(1), pp. 3–23. Available at: <https://doi.org/10.1093/cjres/rsaa004>.
- Collins, C. *et al.* (2021) 'Artificial intelligence in information systems research: A systematic literature review and research agenda', *International Journal of Information Management*, 60, p. 102383. Available at: <https://doi.org/10.1016/j.ijinfomgt.2021.102383>.
- Conner, M. and Norman, P. (2022) 'Understanding the intention-behavior gap: The role of intention strength', *Frontiers in Psychology*, 13, p. 923464. Available at: <https://doi.org/10.3389/fpsyg.2022.923464>.
- Creswell, J.W. (2013) *Qualitative inquiry and research design: choosing among five approaches*. 3rd ed. Los Angeles, Calif: SAGE Publications.
- Crotty, M. (1998) *The foundations of social research: meaning and perspective in the research process*. London ; Thousand Oaks, Calif: Sage Publications.
- Daft, R.L. and Lewin, A.Y. (1990) 'Can Organisation Studies Begin to Break out of the Normal Science Straitjacket?', *Organisation Science*, 1(1), p. pp.1-9.
- Danermark, B., Ekström, M. and Karlsson, J.Ch. (2019) *Explaining society: critical realism in the social sciences*. Second edition. London ; New York: Routledge, Taylor & Francis Group.
- Dastin, J. (2018) 'Amazon scraps secret AI recruiting tool that showed bias against women', *Reuters*, 10 October. Available at: <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G> (Accessed: 21 September 2023).

- Daugherty, P.R. and Wilson, H.J. (2018) *Human + Machine: Reimagining Work in the Age of AI*. Harvard Business Press.
- Davenport, T.H. and Kirby, J. (2016) *Only Humans Need Apply: Winners and Losers in the Age of Smart Machines*. HarperCollins.
- Davidson, E. (2002) 'Technology frames and framing: A Socio-Cognitive investigation of requirements demands', *MIS Quarterly*, 26(4), pp. 329–358. Available at: <https://doi.org/10.2307/4132312>.
- Davidson, E. and Pai, D. (2004) 'Making Sense of Technological Frames: Promise, Progress, and Potential', in B. Kaplan et al. (eds) *Information Systems Research*. Boston, MA: Springer US (IFIP International Federation for Information Processing), pp. 473–491. Available at: https://doi.org/10.1007/1-4020-8095-6_26.
- Davis, F.D. (1989) 'Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology', *MIS Quarterly*, 13(3), p. 319. Available at: <https://doi.org/10.2307/249008>.
- Davis, M.S. (1971) 'That's Interesting!: Towards a Phenomenology of Sociology and a Sociology of Phenomenology', *Philosophy of the Social Sciences*, 1(2), pp. 309–344. Available at: <https://doi.org/10.1177/004839317100100211>.
- Demetis, D. and Lee, A.S. (2018) 'When Humans Using the IT Artifact Becomes IT Using the Human Artifact', *Journal of the Association for Information Systems*, pp. 929–952. Available at: <https://doi.org/10.17705/1jais.00514>.
- Denzin, N.K. and Lincoln, Y.S. (2017) *The SAGE Handbook of Qualitative Research*. Thousand Oaks, Calif: SAGE Publications.
- Department for Business, Energy, and Industrial Strategy, G.B. (2017) *Industrial Strategy: building a Britain fit for the future*. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/664563/industrial-strategy-white-paper-web-ready-version.pdf (Accessed: 21 September 2023).
- Dignam, A. (2020) 'Artificial intelligence, tech corporate governance and the public interest regulatory response', *Cambridge Journal of Regions, Economy and Society*, 13(1), pp. 37–54. Available at: <https://doi.org/10.1093/cjres/rsaa002>.
- DiMaggio, P.J. and Powell, W.W. (1983) 'The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields', *American Sociological Review*, 48(2), pp. 147–160. Available at: <https://doi.org/10.2307/2095101>.
- Du, W. (Derek) and Mao, J.-Y. (2018) 'Developing and maintaining clients' trust through institutional mechanisms in online service markets for digital entrepreneurs: A process model', *The Journal of Strategic Information Systems*, 27(4), pp. 296–310. Available at: <https://doi.org/10.1016/j.jsis.2018.07.001>.
- Easton, G. (2010) 'Critical realism in case study research', *Industrial Marketing Management*, 39(1), pp. 118–128. Available at: <https://doi.org/10.1016/j.indmarman.2008.06.004>.
- Edmondson, A.C. and Mcmanus, S.E. (2007) 'Methodological fit in management field research', *Academy of Management Review*, 32(4), pp. 1246–1264. Available at: <https://doi.org/10.5465/amr.2007.26586086>.

EGOS (no date) *Professions in an Age of Intelligent Technologies - European Group for Organizational Studies*. Available at: https://www.egos.org/jart/prj3/egos/main.jart?rel=de&reserve-mode=active&content-id=1658501434997&subtheme_id=1637466405837&show_prog=yes (Accessed: 22 September 2023).

Eisenhardt, K.M. and Graebner, M.E. (2007) 'Theory Building From Cases: Opportunities And Challenges', *Academy of Management Journal*, 50(1), pp. 25–32. Available at: <https://doi.org/10.5465/amj.2007.24160888>.

Empson, L. *et al.* (eds) (2015) *The Oxford handbook of professional service firms*. First edition. Oxford, United Kingdom: Oxford University Press.

Empson, L. (2017) 'My liminal life: Perpetual journeys across the research-practice divide', in J.M. Bartunek and Mckenzie, J. (eds) *Academic–Practitioner Relationships: Developments, Complexities and Opportunities*. UK: Routledge, p. pp 200-216.

Empson, L. (2018) 'Elite interviewing in professional organizations', *Journal of Professions and Organization*, 5(1), pp. 58–69. Available at: <https://doi.org/10.1093/jpo/jox010>.

Engineering and Physical Sciences Research Council (2022) *Artificial intelligence technologies*. Available at: <https://www.ukri.org/what-we-do/our-main-funds-and-areas-of-support/browse-our-areas-of-investment-and-support/artificial-intelligence-technologies/> (Accessed: 21 September 2023).

Evans, A. and Heimann, A. (2022) *AI Activity in UK Business: An assessment of the scale of AI activity in UK businesses and scenarios for growth over the next twenty years*. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1045381/AI_Activity_in_UK_Businesses_Report__Capital_Economics_and_DCMS__January_2022__Web_accessible_.pdf.

Faggella, D. (2020) *AI in Law and Legal Practice - A Comprehensive View of 35 Current Applications*, *Emerj*. Available at: <https://emerj.com/ai-sector-overviews/ai-in-law-legal-practice-current-applications/> (Accessed: 27 August 2020).

Fairclough, N. (2005) 'Peripheral Vision: Discourse Analysis in Organization Studies: The Case for Critical Realism', *Organization Studies*, 26(6), pp. 915–939. Available at: <https://doi.org/10.1177/0170840605054610>.

Fan, W. *et al.* (2020) 'Investigating the impacting factors for the healthcare professionals to adopt artificial intelligence-based medical diagnosis support system (AIMDSS)', *Annals of Operations Research*, 294(1–2), pp. 567–592. Available at: <https://doi.org/10.1007/s10479-018-2818-y>.

Faraj, S., Pachidi, S. and Sayegh, K. (2018) 'Working and organizing in the age of the learning algorithm', *Information and Organization*, 28(1), pp. 62–70. Available at: <https://doi.org/10.1016/j.infoandorg.2018.02.005>.

Faulconbridge, J., Folke Henriksen, L. and Seabrooke, L. (2021) 'How professional actions connect and protect', *Journal of Professions and Organization*, 8(2), pp. 214–227. Available at: <https://doi.org/10.1093/jpo/joab008>.

Faulconbridge, J. and Muzio, D. (2017) 'Global Professional Service Firms and Institutionalization', in L. Seabrooke and L.F. Henriksen (eds) *Professional Networks in Transnational Governance*.

Cambridge: Cambridge University Press, pp. 219–232. Available at:
<https://doi.org/10.1017/9781316855508.014>.

Faulconbridge, J.R., Sarwar, A. and Spring, M. (2024) 'Accommodating Machine Learning Algorithms in Professional Service Firms', *Organization Studies*, p. 01708406241252930. Available at:
<https://doi.org/10.1177/01708406241252930>.

Fishbein, M. and Ajzen, I. (1975) *Belief, attitude, intention, and behavior: an introduction to theory and research*. Reading, Mass: Addison-Wesley Pub. Co.

Fletcher, M. *et al.* (2018) 'Three Pathways to Case Selection in International Business: A Twenty-Year Review, Analysis and Synthesis', *International Business Review*, 27(4), pp. 755–766. Available at:
<https://doi.org/10.1016/j.ibusrev.2017.12.004>.

Fletcher, M. and Plakoyiannaki, E. (2011) 'Case selection in international business: key issues and common misconceptions', in R. Marschan-Pierkkari and C. Welch (eds). Cheltenham, UK: Edward Elgar, pp. 171–192. Available at: <http://eprints.gla.ac.uk/54991/> (Accessed: 23 September 2023).

Forsythe, D. (2001) *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford University Press.

Frey, C.B. and Osborne, M.A. (2017) 'The future of employment: How susceptible are jobs to computerisation?', *Technological Forecasting and Social Change*, 114, pp. 254–280. Available at:
<https://doi.org/10.1016/j.techfore.2016.08.019>.

Friedland, R. and Alford, R.R. (1991) 'Bringing society back in: Symbols, practices and institutional contradictions.', in *In W. W. Powell & P. DiMaggio (Eds.), The new institutionalism in organizational analysis*. Chicago, IL: University of Chicago Press, p. pp.212-266.

Fryer, T. (2022) 'A critical realist approach to thematic analysis: producing causal explanations', *Journal of Critical Realism*, 21(4), pp. 365–384. Available at:
<https://doi.org/10.1080/14767430.2022.2076776>.

Gefen, D. (2004) 'What Makes an ERP Implementation Relationship Worthwhile: Linking Trust Mechanisms and ERP Usefulness', *Journal of Management Information Systems*, 21(1), pp. 263–288. Available at: <https://doi.org/10.1080/07421222.2004.11045792>.

Gefen, Karahanna, and Straub (2003) 'Trust and TAM in Online Shopping: An Integrated Model', *MIS Quarterly*, 27(1), pp. 51–90. Available at: <https://doi.org/10.2307/30036519>.

George, A.L. and Bennett, A. (2005) *Case Studies and Theory Development in the Social Sciences*. MIT Press.

Glaser, V.L., Sloan, J. and Gehman, J. (2024) 'Organizations as Algorithms: A New Metaphor for Advancing Management Theory', *Journal of Management Studies*, p. joms.13033. Available at:
<https://doi.org/10.1111/joms.13033>.

Goto, M. (2021) 'Collective professional role identity in the age of artificial intelligence', *Journal of Professions and Organization*, 8(1), pp. 86–107. Available at: <https://doi.org/10.1093/jpo/joab003>.

Grace, K. *et al.* (2018) 'Viewpoint: When will AI exceed human performance? Evidence from AI experts', *Journal of Artificial Intelligence Research*, 62. Available at:

<https://ora.ox.ac.uk/objects/uuid:9ef0935d-5ad5-48f7-934b-558f272f50f7> (Accessed: 21 September 2023).

Granić, A. (2023) 'Technology Acceptance and Adoption in Education', in O. Zawacki-Richter and I. Jung (eds) *Handbook of Open, Distance and Digital Education*. Singapore: Springer Nature, pp. 183–197. Available at: https://doi.org/10.1007/978-981-19-2080-6_11.

Gray, L. *et al.* (2020) 'Expanding Qualitative Research Interviewing Strategies: Zoom Video Communications', *The Qualitative Report*, 25(5), pp. 1292–1301. Available at: <https://doi.org/10.46743/2160-3715/2020.4212>.

di Gregorio, S. (2021) 'Voice to text: Automating transcription.', in C. Vanover, P. Mihás, and J. Saldana (eds) *Analyzing and Interpreting Qualitative Research After the Interview*. SAGE Publications Ltd.

Groover, M.P. (2020) *Fundamentals of Modern Manufacturing: Materials, Processes, and Systems*. John Wiley & Sons.

Guest, G., Bunce, A. and Johnson, L. (2006) 'How Many Interviews Are Enough?: An Experiment with Data Saturation and Variability', *Field Methods*, 18(1), pp. 59–82. Available at: <https://doi.org/10.1177/1525822X05279903>.

Ha, N.T. *et al.* (2019) 'The effect of trust on consumers' online purchase intention: An integration of TAM and TPB', *Management Science Letters*, pp. 1451–1460. Available at: <https://doi.org/10.5267/j.msl.2019.5.006>.

Hall, W. and Pesenti, J. (2017) *Growing the artificial intelligence industry in the UK*. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/652097/Growing_the_artificial_intelligence_industry_in_the_UK.pdf (Accessed: 21 September 2023).

Hambrick, D.C. (2007) 'The Field of Management's Devotion to Theory: Too Much of a Good Thing?', *Academy of Management Journal*, 50(6), pp. 1346–1352. Available at: <https://doi.org/10.5465/amj.2007.28166119>.

Harvey (no date) *Harvey | Generative AI for Elite Law Firms*. Available at: <https://www.harvey.ai/> (Accessed: 21 September 2023).

Harvey, J., Heineke, J. and Lewis, M. (2016) 'Editorial for Journal of Operations Management special issue on "Professional Service Operations Management (PSOM)"', *Journal of Operations Management*, 42–43(1), pp. 4–8. Available at: <https://doi.org/10.1016/j.jom.2016.03.005>.

Henfridsson, O. and Bygstad, B. (2013) 'The Generative Mechanisms of Digital Infrastructure Evolution', *MIS Quarterly*, 37(3), pp. 907–931. Available at: <https://doi.org/10.25300/MISQ/2013/37.3.11>.

Hengstler, M., Enkel, E. and Duelli, S. (2016) 'Applied artificial intelligence and trust—The case of autonomous vehicles and medical assistance devices', *Technological Forecasting and Social Change*, 105, pp. 105–120. Available at: <https://doi.org/10.1016/j.techfore.2015.12.014>.

Hennink, M. and Kaiser, B.N. (2022) 'Sample sizes for saturation in qualitative research: A systematic review of empirical tests', *Social Science & Medicine*, 292, p. 114523. Available at: <https://doi.org/10.1016/j.socscimed.2021.114523>.

Hinings, B., Gegenhuber, T. and Greenwood, R. (2018) 'Digital innovation and transformation: An institutional perspective', *Information and Organization*, 28(1), pp. 52–61. Available at: <https://doi.org/10.1016/j.infoandorg.2018.02.004>.

HM Government (2017) 'Industrial Strategy Building a Britain fit for the future'. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/664563/industrial-strategy-white-paper-web-ready-version.pdf (Accessed: 25 August 2020).

Hoddy, E.T. (2019) 'Critical realism in empirical research: employing techniques from grounded theory methodology', *International Journal of Social Research Methodology*, 22(1), pp. 111–124. Available at: <https://doi.org/10.1080/13645579.2018.1503400>.

Hoff, K.A. and Bashir, M. (2015) 'Trust in Automation: Integrating Empirical Evidence on Factors That Influence Trust', *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 57(3), pp. 407–434. Available at: <https://doi.org/10.1177/0018720814547570>.

Holstein, J.A. and Gubrium, J.F. (1995) *The active interview*. Thousand Oaks, CA, US: Sage Publications, Inc (The active interview), pp. vii, 85. Available at: <https://doi.org/10.4135/9781412986120>.

Hongladarom, S. (2020) 'Machine hermeneutics, postphenomenology, and facial recognition technology', *AI & SOCIETY* [Preprint]. Available at: <https://doi.org/10.1007/s00146-020-00951-x>.

Houston, S. (2014) 'Critical Realism', in D. Coghlan and M. Brydon-Miller (eds) *The SAGE Encyclopedia of Action Research*. Thousand Oaks, Calif: SAGE Publications Ltd, pp. 220–222. Available at: <https://doi.org/10.4135/9781446294406>.

Ihde, D. (1990) *Technology and the Lifeworld: From Garden to Earth*. Indiana University Press. Available at: <https://philarchive.org/rec/IHDTAT-3> (Accessed: 23 September 2023).

Institute for Government (2022) *Timeline of UK government coronavirus lockdowns and restrictions*, *Institute for Government*. Available at: <https://www.instituteforgovernment.org.uk/data-visualisation/timeline-coronavirus-lockdowns> (Accessed: 23 September 2023).

Jarvenpaa, S.L., Knoll, K. and Leidner, D.E. (1998) 'Is Anybody out There? Antecedents of Trust in Global Virtual Teams', *Journal of Management Information Systems*, 14(4), pp. 29–64.

Jenner, B.M. and Myers, K.C. (2019) 'Intimacy, rapport, and exceptional disclosure: a comparison of in-person and mediated interview contexts', *International Journal of Social Research Methodology*, 22(2), pp. 165–177. Available at: <https://doi.org/10.1080/13645579.2018.1512694>.

Jensen, T.B. and Aanestad, M. (2007) 'Hospitality and hostility in hospitals: a case study of an EPR adoption among surgeons', *European Journal of Information Systems*, 16(6), pp. 672–680. Available at: <https://doi.org/10.1057/palgrave.ejis.3000713>.

Jussupow, E., Heinzl, A. and Spohrer, K. (2018) 'I am; We are - Conceptualizing Professional Identity Threats from Information Technology', in. *Thirty Ninth International Conference on Information Systems, San Francisco 2018*, p. 17.

Karjalainen, M., Niemistö, C. and Hearn, J. (2015) 'Unpacking the problem of research access(es): The case of large knowledge-intensive international consultancies', *Qualitative Research in Organizations and Management: An International Journal*, 10(3), pp. 274–293. Available at: <https://doi.org/10.1108/QROM-12-2013-1189>.

- Karunamuni, N., Imayama, I. and Goonetilleke, D. (2021) 'Pathways to well-being: Untangling the causal relationships among biopsychosocial variables', *Social Science & Medicine* (1982), 272, p. 112846. Available at: <https://doi.org/10.1016/j.socscimed.2020.112846>.
- Katz, D.M. (2013) 'Quantitative Legal Prediction – or – How I Learned to Stop Worrying and Start Preparing for the Data Driven Future of the Legal Services Industry', *Emory Law Journal*, 62, pp. 909–966.
- Kelemen, M. and Rumens, N. (2008) *An introduction to critical management research*. Los Angeles ; London: SAGE.
- Kellogg, K.C. (2022) 'Local Adaptation Without Work Intensification: Experimentalist Governance of Digital Technology for Mutually Beneficial Role Reconfiguration in Organizations', *Organization Science*, 33(2), pp. 571–599. Available at: <https://doi.org/10.1287/orsc.2021.1445>.
- Kempster, S. (2006) 'Leadership learning through lived experience: A process of apprenticeship?', *Journal of Management & Organization*, 12(1), pp. 4–22. Available at: <https://doi.org/10.5172/jmo.2006.12.1.4>.
- Kempster, S. and Parry, K.W. (2011) 'Grounded theory and leadership research: A critical realist perspective', *The Leadership Quarterly*, 22(1), pp. 106–120. Available at: <https://doi.org/10.1016/j.leaqua.2010.12.010>.
- Kira Systems (no date) *Machine Learning Contract Search, Review and Analysis Software*. Available at: <https://kirasystems.com/> (Accessed: 21 September 2023).
- Köktener, B. and Tunçalp, D. (2021) 'Old game, new rules and “odd friends”: Digitalization, jurisdictional conflicts, and boundary work of auditors in a “big four” professional service firm', *Journal of Professions and Organization*, 8(3), pp. 349–373. Available at: <https://doi.org/10.1093/jpo/joab016>.
- Korica, M. and Molloy, E. (2010) 'Making sense of professional identities: Stories of medical professionals and new technologies', *Human Relations*, 63(12), pp. 1879–1901. Available at: <https://doi.org/10.1177/0018726710367441>.
- Kowal, S. and O'Connell, D.C. (2014) 'Transcription as a Crucial Step of Data Analysis', in U. Flick (ed.) *The SAGE Handbook of Qualitative Data Analysis*. SAGE Publications Ltd. Available at: <https://doi.org/10.4135/9781446282243>.
- von Krogh, G. (2018) 'Artificial Intelligence in Organizations: New Opportunities for Phenomenon-Based Theorizing', *Academy of Management Discoveries*, 4(4), pp. 404–409. Available at: <https://doi.org/10.5465/amd.2018.0084>.
- von Krogh, G., Rossi-Lamastra, C. and Haefliger, S. (2012) 'Phenomenon-based Research in Management and Organisation Science: When is it Rigorous and Does it Matter?', *Long Range Planning*, 45(4), pp. 277–298. Available at: <https://doi.org/10.1016/j.lrp.2012.05.001>.
- Kronblad, C. (2020) 'How Digitalization Changes our Understanding of Professional Service Firms', *Academy of Management Discoveries*, 6(3), pp. 436–454. Available at: <https://doi.org/10.5465/amd.2019.0027>.

- Lankton, N., McKnight, D.H. and Thatcher, J.B. (2014) 'Incorporating trust-in-technology into Expectation Disconfirmation Theory', *The Journal of Strategic Information Systems*, 23(2), pp. 128–145. Available at: <https://doi.org/10.1016/j.jsis.2013.09.001>.
- Lankton, N., McKnight, D.H. and Tripp, J. (2015) 'Technology, Humanness, and Trust: Rethinking Trust in Technology', *Journal of the Association for Information Systems*, 16(10), pp. 880–918. Available at: <https://doi.org/10.17705/1jais.00411>.
- Lee, A.S., Thomas, M. and Baskerville, R.L. (2015) 'Going back to basics in design science: from the information technology artifact to the information systems artifact', *Information Systems Journal*, 25(1), pp. 5–21. Available at: <https://doi.org/10.1111/isj.12054>.
- Lee, J.D. and See, K.A. (2004) 'Trust in Automation: Designing for Appropriate Reliance', *Human Factors*, 46(1), pp. 50–80. Available at: https://doi.org/10.1518/hfes.46.1.50_30392.
- Lee, R.S.T. (2020) *Artificial Intelligence in Daily Life*. 1st ed. 2020 edition. Singapore: Springer.
- Leppäaho, T., Plakoyiannaki, E. and Dimitratos, P. (2016) 'The Case Study in Family Business: An Analysis of Current Research Practices and Recommendations', *Family Business Review*, 29(2), pp. 159–173. Available at: <https://doi.org/10.1177/0894486515614157>.
- Li, X., Hess, T.J. and Valacich, J.S. (2008) 'Why do we trust new technology? A study of initial trust formation with organizational information systems', *The Journal of Strategic Information Systems*, 17(1), pp. 39–71. Available at: <https://doi.org/10.1016/j.jsis.2008.01.001>.
- Lim, S. et al. (2013) 'Theories Used in Information Systems Research: Insights from Complex Network Analysis', *Journal of Information Technology Theory & Application*, 14(2), pp. 5–46.
- Lindgren, S. and Holmström, J. (2020) 'A Social Science Perspective on Artificial Intelligence: Building Blocks for a Research Agenda', *Journal of Digital Social Research*, 2(3), pp. 1–15. Available at: <https://doi.org/10.33621/jdsr.v2i3.65>.
- Liu, B.Q. and Goodhue, D.L. (2012) 'Two Worlds of Trust for Potential E-Commerce Users: Humans as Cognitive Misers', *Information Systems Research*, 23(4), pp. 1246–1262. Available at: <https://doi.org/10.1287/isre.1120.0424>.
- Lobe, B., Morgan, D. and Hoffman, K.A. (2020) 'Qualitative Data Collection in an Era of Social Distancing', *International Journal of Qualitative Methods*, 19, p. 1609406920937875. Available at: <https://doi.org/10.1177/1609406920937875>.
- Looker, B., Vickers, J. and Kington, A. (2021) 'The Case for a Critical Realist Grounded Theory Research Design', in I. Psaroudakis, T. Muller, and A. Salvini (eds). Pisa: Pisa University Press, pp. 139–168. Available at: https://doi.org/10.1/Dealing_With_Grounded_Theory_%20looker%20vickers%20kington.pdf.
- Lumineau, F., Schilke, O. and Wang, W. (2023) 'Organizational Trust in the Age of the Fourth Industrial Revolution: Shifts in the Form, Production, and Targets of Trust', *Journal of Management Inquiry*, 32(1), pp. 21–34. Available at: <https://doi.org/10.1177/10564926221127852>.
- Malone, T., Rus, D. and Laubacher, R. (2020) *Artificial intelligence and the future of work*. Available at: <https://workofthefuture.mit.edu/research-post/artificial-intelligence-and-the-future-of-work/>.

- Markou, C. and Deakin, S. (2020) 'Is Law Computable? From Rule of Law to Legal Singularity'. Rochester, NY. Available at: <https://doi.org/10.2139/ssrn.3589184>.
- Mayer, R.C., Davis, J.H. and Schoorman, F.D. (1995) 'An Integrative Model of Organizational Trust', *The Academy of Management Review*, 20(3), pp. 709–734. Available at: <https://doi.org/10.2307/258792>.
- McCulloch, W.S. and Pitts, W. (1943) 'A logical calculus of the ideas immanent in nervous activity', *The bulletin of mathematical biophysics*, 5(4), pp. 115–133. Available at: <https://doi.org/10.1007/BF02478259>.
- McGinnis, J.O. and Pearce, R.G. (2014) 'The Great Disruption: How Machine Intelligence Will Transform the Role of Lawyers in the Delivery of Legal Services', *Fordham Legal Review*, 82(6), pp. 3041–3066. Available at: <https://doi.org/10.21202/1993-047X.13.2019.2.1230-1250>.
- Mcknight, D.H. et al. (2011) 'Trust in a specific technology: An investigation of its components and measures', *ACM Transactions on Management Information Systems*, 2(2), pp. 1–25. Available at: <https://doi.org/10.1145/1985347.1985353>.
- McKnight, D.H., Choudhury, V. and Kacmar, C. (2002) 'The impact of initial consumer trust on intentions to transact with a web site: a trust building model', *The Journal of Strategic Information Systems*, 11(3–4), pp. 297–323. Available at: [https://doi.org/10.1016/S0963-8687\(02\)00020-3](https://doi.org/10.1016/S0963-8687(02)00020-3).
- McLuhan, M. and McLuhan, E. (1988) *Laws of Media: The New Science*. University of Toronto Press.
- McMullin, C. (2023) 'Transcription and Qualitative Methods: Implications for Third Sector Research', *Voluntas: International Journal of Voluntary and Nonprofit Organizations*, 34(1), pp. 140–153. Available at: <https://doi.org/10.1007/s11266-021-00400-3>.
- Milmo, D. (2023) 'ChatGPT reaches 100 million users two months after launch', *The Guardian*, 2 February. Available at: <https://www.theguardian.com/technology/2023/feb/02/chatgpt-100-million-users-open-ai-fastest-growing-app> (Accessed: 21 September 2023).
- Mingers, J., Mutch, A. and Willcocks, L. (2013) 'Critical realism in information systems research', *MIS Quarterly*, 37(3), pp. 795–802. Available at: <https://doi.org/10.25300/MISQ/2013/37:3.3>.
- Mintzberg, H. (1978) 'Patterns in Strategy Formation', *Management Science*, 24(9), pp. 934–948. Available at: <https://doi.org/10.1287/mnsc.24.9.934>.
- Myers and Klein (2011) 'A Set of Principles for Conducting Critical Research in Information Systems', *MIS Quarterly*, 35(1), p. 17. Available at: <https://doi.org/10.2307/23043487>.
- Nelson, A.J. and Irwin, J. (2014) "'Defining What We Do—All Over Again": Occupational Identity, Technological Change, and the Librarian/Internet-Search Relationship', *Academy of Management Journal*, 57(3), pp. 892–928. Available at: <https://doi.org/10.5465/amj.2012.0201>.
- Nilsson, N.J. (1971) *Problem-solving Methods in Artificial Intelligence*. McGraw-Hill.
- Noordegraaf, M. (2015) 'Hybrid professionalism and beyond: (New) Forms of public professionalism in changing organizational and societal contexts', *Journal of Professions and Organization*, 2(2), pp. 187–206. Available at: <https://doi.org/10.1093/jpo/jov002>.

Noordegraaf, M. (2020) 'Protective or connective professionalism? How connected professionals can (still) act as autonomous and authoritative experts', *Journal of Professions and Organization*, 7(2), pp. 205–223. Available at: <https://doi.org/10.1093/jpo/joaa011>.

von Nordenflycht, A. (2010) 'What Is a Professional Service Firm? Toward a Theory and Taxonomy of Knowledge-Intensive Firms', *Academy of Management Review*, 35(1), pp. 155–174. Available at: <https://doi.org/10.5465/amr.35.1.zok155>.

von Nordenflycht, A. (2014) 'Does the emergence of publicly traded professional service firms undermine the theory of the professional partnership? A cross-industry historical analysis', *Journal of Professions and Organization*, 1(2), pp. 137–160. Available at: <https://doi.org/10.1093/jpo/jou002>.

Norman, D.A. and Draper, S.W. (eds) (1986) *User centered system design: new perspectives on human-computer interaction*. Hillsdale, N.J.: Erlbaum.

Odendahl, T. and Shaw, A. (2002) 'Interviewing elites', in J.F. Gubrium and J.A. Holstein (eds) *Handbook of Interview Research: Context and Method*. SAGE.

Okhuysen, G.A. and Bechky, B.A. (2009) '10 Coordination in Organizations: An Integrative Perspective', *Academy of Management Annals*, 3(1), pp. 463–502. Available at: <https://doi.org/10.5465/19416520903047533>.

Oliver, C. (2012) 'Critical Realist Grounded Theory: A New Approach for Social Work Research', *The British Journal of Social Work*, 42(2), pp. 371–387. Available at: <https://doi.org/10.1093/bjsw/bcr064>.

Orlikowski, W.J. and Baroudi, J.J. (1991) 'Studying Information Technology in Organizations: Research Approaches and Assumptions', *Information Systems Research*, 2(1), pp. 1–28.

Orlikowski, W.J. and Gash, D.C. (1994) 'Technological frames: making sense of information technology in organizations', *ACM Transactions on Information Systems*, 12(2), pp. 174–207. Available at: <https://doi.org/10.1145/196734.196745>.

Orlikowski, W.J. and Scott, S.V. (2008) 'Sociomateriality: Challenging the separation of technology, work and organization', *The Academy of Management Annals*, 2(1), pp. 433–474. Available at: <https://doi.org/10.1080/19416520802211644>.

Pareliussen, B., Æsøy, V. and Giskeødegård, M.F. (2022) 'Professions, work, and digitalization: Technology as means to connective professionalism', *Journal of Professions and Organization*, 9(1), pp. 100–114. Available at: <https://doi.org/10.1093/jpo/joab023>.

Patel, B.N. *et al.* (2019) 'Human–machine partnership with artificial intelligence for chest radiograph diagnosis', *npj Digital Medicine*, 2(1), pp. 1–10. Available at: <https://doi.org/10.1038/s41746-019-0189-7>.

Paton, S. and Hodgson, D. (2016) 'Project managers on the edge: liminality and identity in the management of technical work', *New Technology, Work and Employment*, 31(1), pp. 26–40. Available at: <https://doi.org/10.1111/ntwe.12056>.

Patton, M.Q. (2014) *Qualitative Research & Evaluation Methods: Integrating Theory and Practice*. SAGE Publications.

- Paulus, T., Lester, J. and Dempster, P. (2013) *Digital Tools for Qualitative Research*. SAGE.
- Paulus, T.M. and Lester, J.N. (2021) *Doing Qualitative Research in a Digital World*. Thousand Oaks, Calif: SAGE.
- Pawson, R. and Tilley, N. (1997) *Realistic evaluation*. London ; Thousand Oaks, Calif: Sage.
- Petintseva, O., Faria, R. and Eski, Y. (2019) *Interviewing Elites, Experts and the Powerful in Criminology*. Springer Nature.
- Petriglieri, J.L. (2011) 'Under Threat: Responses to and the Consequences of Threats to Individuals' Identities', *Academy of Management Review*, 36(4), pp. 641–662. Available at: <https://doi.org/10.5465/amr.2009.0087>.
- Piekkari, R., Welch, C. and Paavilainen, E. (2009) 'The Case Study as Disciplinary Convention: Evidence From International Business Journals', *Organizational Research Methods*, 12(3), pp. 567–589. Available at: <https://doi.org/10.1177/1094428108319905>.
- Pilgrim, D. (2020) *Critical Realism for Psychologists*. 1st edn. London ; New York: Routledge. Available at: <https://doi.org/10.4324/9780429274497>.
- Pine, K.H. and Mazmanian, M. (2017) 'Artful and Contorted Coordinating: The Ramifications of Imposing Formal Logics of Task Jurisdiction on Situated Practice', *Academy of Management Journal*, 60(2), pp. 720–742. Available at: <https://doi.org/10.5465/amj.2014.0315>.
- Plakoyiannaki, E. (2020) 'Qualitative Case Study Research'. Available at: <https://narti.org.uk/event/qualitative-case-study-research-with-professor-emmanuella-plakoyiannaki-online/>.
- Point, S. and Baruch, Y. (2023) '(Re)thinking transcription strategies: Current challenges and future research directions', *Scandinavian Journal of Management*, 39(2), p. 101272. Available at: <https://doi.org/10.1016/j.scaman.2023.101272>.
- Rahman, S.A. et al. (2021) 'Resilient Research in the Field: Insights and Lessons From Adapting Qualitative Research Projects During the COVID-19 Pandemic', *International Journal of Qualitative Methods*, 20, p. 16094069211016106. Available at: <https://doi.org/10.1177/16094069211016106>.
- Rai, A., Constantinides, P. and Sarker, S. (2019) 'Next-Generation Digital Platforms: Toward Human-AI hybrids', *MIS Quarterly*, 43(1), pp. iii–ix.
- Raisch, S. and Krakowski, S. (2021) 'Artificial intelligence and management: The automation–augmentation paradox', *The Academy of Management Review*, 46(1), pp. 192–210. Available at: <https://doi.org/10.5465/amr.2018.0072>.
- Ramalho, R. et al. (2015) 'Literature Review and Constructivist Grounded Theory Methodology', *Forum Qualitative Sozialforschung / Forum: Qualitative Social Research*, 16(3). Available at: <https://doi.org/10.17169/fqs-16.3.2313>.
- Reay, T. et al. (2017) 'Getting Leopards to Change their Spots: Co-creating a New Professional Role Identity', *Academy of Management Journal*, 60(3), pp. 1043–1070. Available at: <https://doi.org/10.5465/amj.2014.0802>.
- Richard, D.L. (2002) 'Herding Cats: The Lawyer Personality Revealed', *Altman Weil*, 29(11), pp. 2–12.

- Richardson, J., Godfrey, B. and Walklate, S. (2021) 'Rapid, remote and responsive research during COVID-19', *Methodological Innovations*, 14(1), p. 20597991211008581. Available at: <https://doi.org/10.1177/20597991211008581>.
- Rose, J. and Schlichter, B.R. (2013) 'Decoupling, re-engaging: managing trust relationships in implementation projects: Managing trust relationships in implementation projects', *Information Systems Journal*, 23(1), pp. 5–33. Available at: <https://doi.org/10.1111/j.1365-2575.2011.00392.x>.
- Rousseau, D.M. *et al.* (1998) 'Introduction to Special Topic Forum: Not so Different after All: A Cross-Discipline View of Trust', *The Academy of Management Review*, 23(3), pp. 393–404.
- Russell, S. and Norvig, P. (2016) *Artificial Intelligence: A Modern Approach*. CreateSpace Independent Publishing Platform.
- Ryan, M. (2020) 'In AI We Trust: Ethics, Artificial Intelligence, and Reliability', *Science and Engineering Ethics*, 26(5), pp. 2749–2767. Available at: <https://doi.org/10.1007/s11948-020-00228-y>.
- Sako, M. and Parnham, R. (2021) 'Technology and Innovation in Legal Services: Final Report for the Solicitors Regulation Authority.' Available at: <https://www.sra.org.uk/globalassets/documents/sra/research/full-report-technology-and-innovation-in-legal-services.pdf?version=4a1bfe>.
- Sako, M., Qian, M. and Attolini, J. (2022) 'Future of professional work: evidence from legal jobs in Britain and the United States', *Journal of Professions and Organization*, 9(2), pp. 143–169. Available at: <https://doi.org/10.1093/jpo/joac011>.
- Saldana, J. (2016) *The Coding Manual for Qualitative Researchers*. SAGE.
- Salmons, J. (2014) *Qualitative Online Interviews: Strategies, Design, and Skills*. Thousand Oaks, Calif: SAGE Publications.
- Sandelowski, M. (1995) 'Qualitative analysis: what it is and how to begin', *Research in Nursing & Health*, 18(4), pp. 371–375. Available at: <https://doi.org/10.1002/nur.4770180411>.
- Saunders, M.N.K., Lewis, P. and Thornhill, A. (2015) *Research methods for business students*. Seventh edition. New York: Pearson Education.
- Saunders, M.N.K. and Townsend, K. (2016) 'Reporting and Justifying the Number of Interview Participants in Organization and Workplace Research', *British Journal of Management*, 27(4), pp. 836–852. Available at: <https://doi.org/10.1111/1467-8551.12182>.
- Schaefer, K.E. *et al.* (2016) 'A Meta-Analysis of Factors Influencing the Development of Trust in Automation: Implications for Understanding Autonomy in Future Systems', *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 58(3), pp. 377–400. Available at: <https://doi.org/10.1177/0018720816634228>.
- Schlichter, B.R. and Rose, J. (2013) 'Trust dynamics in a large system implementation: six theoretical propositions', *European Journal of Information Systems*, 22(4), pp. 455–474. Available at: <https://doi.org/10.1057/ejis.2012.24>.
- Schuetz, S. and Venkatesh, V. (2020) 'Research Perspectives: The Rise of Human Machines: How Cognitive Computing Systems Challenge Assumptions of User-System Interaction', *Journal of the*

- Association for Information Systems*, pp. 460–482. Available at: <https://doi.org/10.17705/1jais.00608>.
- Sherer, P.D. (2019) 'The researcher as detective: Another way of looking at professions and organizations', *Journal of Professions and Organization*, 6(1), pp. 91–102. Available at: <https://doi.org/10.1093/jpo/joy020>.
- Silverman, D. (2020) 'Collecting qualitative data during a pandemic.', *Communication and Medicine*, 17(1), pp. 76–84.
- Smets, M. et al. (2017) '25 years since "P2": Taking stock and charting the future of professional firms', *Journal of Professions and Organization*, 4(2), pp. 91–111. Available at: <https://doi.org/10.1093/jpo/jox006>.
- Smith, M.L. (2006) 'Overcoming theory-practice inconsistencies: Critical realism and information systems research', *Information and Organization*, 16(3), pp. 191–211. Available at: <https://doi.org/10.1016/j.infoandorg.2005.10.003>.
- Söllner, M., Hoffmann, A. and Leimeister, J.M. (2016) 'Why different trust relationships matter for information systems users', *European Journal of Information Systems*, 25(3), pp. 274–287. Available at: <https://doi.org/10.1057/ejis.2015.17>.
- Solomonic* (no date) *Solomonic*. Available at: <https://www.solomonic.co.uk> (Accessed: 21 September 2023).
- Spiehl, P. et al. (2021) 'Technological Frames in the Digital Age: Theory, Measurement Instrument, and Future Research Areas', *Journal of Management Studies*, 58(7), pp. 1962–1993. Available at: <https://doi.org/10.1111/joms.12720>.
- Spring, M., Faulconbridge, J. and Sarwar, A. (2022) 'How information technology automates and augments processes: Insights from Artificial-Intelligence-based systems in professional service operations', *Journal of Operations Management*, 68(6–7), pp. 592–618. Available at: <https://doi.org/10.1002/joom.1215>.
- Stephens, N. (2007) 'Collecting data from elites and ultra elites: telephone and face-to-face interviews with macroeconomists', *Qualitative Research*, 7(2), pp. 203–216. Available at: <https://doi.org/10.1177/1468794107076020>.
- Suddaby, R. and Greenwood, R. (2005) 'Rhetorical Strategies of Legitimacy', *Administrative Science Quarterly*, 50(1), pp. 35–67. Available at: <https://doi.org/10.2189/asqu.2005.50.1.35>.
- Suddaby, R., Greenwood, R. and Wilderom, C. (2008) 'Introduction to the Journal of Organizational Behavior's special issue on professional service firms: where organization theory and organizational behavior might meet', *Journal of Organizational Behavior*, 29(8), pp. 989–994. Available at: <https://doi.org/10.1002/job.552>.
- Susskind, R.E. and Susskind, D. (2015) *The future of the professions: how technology will transform the work of human experts*. First edition. Oxford, United Kingdom: Oxford University Press.
- Symon, G. and Cassell, C. (2012) *Qualitative Organizational Research: Core Methods and Current Challenges*. 55 City Road, London. Available at: <https://doi.org/10.4135/9781526435620>.

Tao, D. (2009) 'Intention to use and actual use of electronic information resources: further exploring Technology Acceptance Model (TAM)', *AMIA ... Annual Symposium proceedings. AMIA Symposium*, 2009, pp. 629–633.

The Law Society (2019) *Lawtech adoption research*. Available at: <https://www.lawsociety.org.uk/topics/research/lawtech-adoption-report> (Accessed: 21 September 2023).

Thompson, T.L. and Graham, B. (2021) 'A more-than-human approach to researching AI at work: Alternative narratives for human and AI co-workers.', in *Bonderup Dohn N, Jørgen Hansen J, Børsen Hansen S, Ryberg T & de Laat M (eds.) Conceptualizing and innovating education and work with networked learning*. Cham, Switzerland: Springer International Publishing, pp. pp171-187.

Thurlow, L. (2020) 'Grounded theory and the PhD – notes for novice researchers', *Journal of Humanities and Applied Social Sciences*, 2(4), pp. 257–270. Available at: <https://doi.org/10.1108/JHASS-05-2020-0079>.

Tredinnick, L. (2017) 'Artificial intelligence and professional roles', *Business Information Review*, 34(1), pp. 37–41. Available at: <https://doi.org/10.1177/0266382117692621>.

Tremblay, S. *et al.* (2021) 'Conducting Qualitative Research to Respond to COVID-19 Challenges: Reflections for the Present and Beyond', *International Journal of Qualitative Methods*, 20, p. 16094069211009679. Available at: <https://doi.org/10.1177/16094069211009679>.

Turing, A.M. (1950) 'I.—COMPUTING MACHINERY AND INTELLIGENCE', *Mind*, LIX(236), pp. 433–460. Available at: <https://doi.org/10.1093/mind/LIX.236.433>.

Turner, V. and Abrahams, R.D. (2017) *The Ritual Process: Structure and Anti-Structure*. 1st edn. Routledge. Available at: <https://doi.org/10.4324/9781315134666>.

UK AI Council (2021) *AI Roadmap*, GOV.UK. Available at: <https://www.gov.uk/government/publications/ai-roadmap> (Accessed: 21 September 2023).

UK Parliament (2018) *AI in the UK: ready, willing and able*, pp. 1–183. Available at: <https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf>.

UKRI (no date) *Next generation services*. Available at: <https://www.ukri.org/what-we-do/our-main-funds-and-areas-of-support/browse-our-areas-of-investment-and-support/next-generation-services/> (Accessed: 22 September 2023).

University of Sheffield (no date) *University's Ethics Policy Governing Research Involving Human Participants, Personal Data and Human Tissue Version 7.5*. Available at: <https://www.sheffield.ac.uk/rpi/ethics-integrity/policy> (Accessed: 23 September 2023).

Van Audenhove, L. and Donders, K. (2019) 'Talking to People III: Expert Interviews and Elite Interviews', in H.V. den Bulck *et al.* (eds) *The Palgrave Handbook of Methods for Media Policy Research*. Springer.

Van de Ven, A.H. (2007) *Engaged scholarship: a guide for organizational and social research*. Oxford ; New York: Oxford University Press.

- Van De Ven, A.H. and Johnson, P.E. (2006) 'Knowledge for Theory and Practice', *Academy of Management Review*, 31(4), pp. 802–821. Available at: <https://doi.org/10.5465/amr.2006.22527385>.
- Vance, A., Elie-Dit-Cosaque, C. and Straub, D.W. (2008) 'Examining Trust in Information Technology Artifacts: The Effects of System Quality and Culture', *Journal of Management Information Systems*, 24(4), pp. 73–100. Available at: <https://doi.org/10.2753/MIS0742-1222240403>.
- Venkatesh *et al.* (2003) 'User Acceptance of Information Technology: Toward a Unified View', *MIS Quarterly*, 27(3), pp. 425–478. Available at: <https://doi.org/10.2307/30036540>.
- Venkatesh, V., Brown, S.A. and Bala, H. (2013) 'Bridging the Qualitative-Quantitative Divide: Guidelines for Conducting Mixed Methods Research in Information Systems', *MIS Quarterly*, 37(1), pp. 21–54. Available at: <https://doi.org/10.25300/MISQ/2013/37.1.02>.
- Venkatesh, V. and Davis, F.D. (1996) 'A Model of the Antecedents of Perceived Ease of Use: Development and Test*', *Decision Sciences*, 27(3), pp. 451–481. Available at: <https://doi.org/10.1111/j.1540-5915.1996.tb00860.x>.
- Venkatesh, V. and Davis, F.D. (2000) 'A Theoretical Extension of the Technology Acceptance Model: Four Longitudinal Field Studies', *Management Science*, 46(2), pp. 186–204.
- Vincent, S. and O'Mahoney, J. (2018) 'Critical Realism and Qualitative Research: An Introductory Overview', in Cassell, C., Cunliffe, A., and Grandy, G., *The SAGE Handbook of Qualitative Business and Management Research Methods: History and Traditions*. London, United Kingdom: SAGE Publications Ltd, pp. 201–216. Available at: <https://doi.org/10.4135/9781526430212.n13>.
- Volkoff, O. and Strong, D.M. (2013) 'Critical Realism and Affordances: Theorizing IT-Associated Organizational Change Processes', *MIS Quarterly*, 37(3), pp. 819–834. Available at: <https://doi.org/10.25300/MISQ/2013/37.3.07>.
- Walsham, G. (1995) 'Interpretive case studies in IS research: nature and method', *European Journal of Information Systems*, 4, pp. 74–81.
- Wang, W. and Benbasat, I. (2008) 'Attributions of Trust in Decision Support Technologies: A Study of Recommendation Agents for E-Commerce', *Journal of Management Information Systems*, 24(4), pp. 249–273. Available at: <https://doi.org/10.2753/MIS0742-1222240410>.
- Wang, S. and Liang, Z. (2024) 'What does the public think about artificial intelligence? An investigation of technological frames in different technological context', *Government Information Quarterly*, 41(2), p. 101939. Available at: <https://doi.org/10.1016/j.giq.2024.101939>.
- Waxman, I. (2022) *AI vs Lawyers - The Ultimate Showdown, Superlegal*. Available at: <https://www.superlegal.ai/blog/aivslawyer/> (Accessed: 21 September 2023).
- Welch, C. *et al.* (2011) 'Theorising from case studies: Towards a pluralist future for international business research', *Journal of International Business Studies*, 42(5), pp. 740–762.
- Welch, C. *et al.* (2022) 'Reconciling theory and context: How the case study can set a new agenda for international business research', *Journal of International Business Studies*, 53(1), pp. 4–26.
- Wellner, G. (2020) 'Material Hermeneutic of Digital Technologies in the Age of AI', *AI and Society*, pp. 1–8. Available at: <https://doi.org/10.1007/s00146-020-00952-w>.

Williams, C.K. and Wynn, D.E. (2018) 'A critical realist script for creative theorising in information systems', *European Journal of Information Systems*. Edited by N.R. Hassan, J. Mingers, and B. Stahl, 27(3), pp. 315–325. Available at: <https://doi.org/10.1080/0960085X.2018.1435231>.

Wiltshire, G. and Ronkainen, N. (2021) 'A realist approach to thematic analysis: making sense of qualitative data through experiential, inferential and dispositional themes', *Journal of Critical Realism*, 20(2), pp. 159–180. Available at: <https://doi.org/10.1080/14767430.2021.1894909>.

Wright, S. (2023) *Creating and Working With Auto-Transcripts in NVivo*. SAGE Publications Ltd. Available at: <https://doi.org/10.4135/9781529629170>.

Wynn, D.E. and Williams, C.K. (2020) 'Recent Advances and Opportunities for Improving Critical Realism-Based Case Study Research in IS', *Journal of the Association for Information Systems*, pp. 50–89. Available at: <https://doi.org/10.17705/1jais.00592>.

Wynn and Williams (2012) 'Principles for Conducting Critical Realist Case Study Research in Information Systems', *MIS Quarterly*, 36(3), p. 787. Available at: <https://doi.org/10.2307/41703481>.

Yin, R.K. (2018) *Case study research and applications: design and methods*. Sixth edition. Los Angeles: SAGE.

Zachariadis, M., Scott, S.V. and Barrett, M. (2013) 'Methodological Implications of Critical Realism for Mixed-Methods Research', *MIS Quarterly*, 37(3), pp. 855–879. Available at: <https://doi.org/10.25300/MISQ/2013/37.3.09>.

Zheng, J. (Frank) and Jarvenpaa, S.L. (2021) 'Thinking Technology as Human: Affordances, Technology Features, and Egocentric Biases in Technology Anthropomorphism', *Journal of the Association for Information Systems*, 22(5), pp. 1429–1453. Available at: <https://doi.org/10.17705/1jais.00698>.

Zion Market Research (2019) *Global LegalTech Artificial Intelligence Market Worth Reach USD 37,858 Million By 2026*. Available at: <https://www.zionmarketresearch.com/news/legaltech-artificial-intelligence-market> (Accessed: 27 August 2020).

Zoom Video Communications Inc (no date) *Zoom, Zoom*. Available at: <https://zoom.us/> (Accessed: 23 September 2023).

Zuckerman, H. (1972) 'Interviewing an ultra-elite', *Public Opinion Quarterly*, 36(2), pp. 159–175. Available at: <https://doi.org/10.1086/267989>.

APPENDICES

Appendix 1: Interview guide

[NB: ensure technical aspects of Zoom are working properly + advise about turning off self-view; start recording the interview]

Introduction

Thank you for agreeing to take part in this research project. I want to encourage you to take an active role in the research process, generating data and insights in partnership with me. Today's interview will differ somewhat from those you may be more familiar with, for example interviewing for a new role or watching a journalist interview a politician. In those interviews there is often a stark imbalance of power between those involved; 'correct' answers that are being sought; and an element of 'performance' whereby those involved try to present themselves in a positive light.

As far as it is possible to do so, please try to relax and enjoy our discussion today. My aim is to give you an opportunity to share your insights about AI in the context of your work. I will ask you a small number of open-ended questions, that you can choose how to answer. I may also on occasion ask follow-up questions based on your initial response.

The data we generate between us will not be directly attributable to you and please do not feel under any pressure to answer questions, when you either have nothing to say, or you wish to keep your views private.

Before we begin the interview can you confirm that you have given your informed consent to participate in this interview based on the information that was previously provided to you in the Information Sheet for this research project.

[capture consent]

Do you have any questions you would like to ask me before we begin?

Building rapport

Before asking you any questions I wanted to provide you with a bit of background about myself and how I came to be researching this topic.

I've now spent over a decade working in the professional service sector. First in the HR department of a commercial law firm, where I was responsible for recruiting and developing junior lawyers. I now work as consultant to law firms and in-house legal teams and coach lawyers at various stages of their career.

Alongside my consultancy work I also retrained as an organisational psychologist before commencing my PhD, which focuses on how lawyers experience the implementation of artificial intelligence at work. How professionals respond to technology has been a long-standing interest of mine. Growing up I remember my father, who was a lawyer, struggling to adapt to the digitalisation of his profession and more recently I have witnessed the growth of the LegalTech industry and the introduction of new technologies, such as those that we will discuss today.

Today's interview sees me adopt my identity as a researcher rather than consultant. The aim of the research is, therefore, to enhance our theoretical understanding of AI as a phenomenon, it is not a project commissioned by the firm to help it meet its own commercial goals. However, as with all academic research, I hope that any insights the research generates will bring practical benefits to professionals working with AI.

For transparency, prior to today's interview I have conducted some general information gathering at the firm about the use of AI, which also included demonstrations of the different technologies being developed or which are already in use. The aim of this was to help orientate me ahead of our discussions, so I would be aware of some of the terminology and technologies you might refer to. This should mean we can focus in this interview on your experiences of using AI.

Questions

Section 1 – General background

I want to start by collecting a bit of biographical information about yourself. Please describe to me your current role at the firm.

I appreciate you referenced some of this in the reflective exercise, but what experience do you have of using LegalTech and specifically AI in your current and previous roles?

Possible probes: Which AI do you have personal experience using?

Tell me more about some of the specific practices linked to using these technologies? Is anyone/thing else involved?

AI can be a bit of a slippery term, how do you define AI in the context of your own work?

Possible probes: What is it about the technology you use that makes them an example of AI?

How did you arrive at this definition – what sources are you drawing upon?

Do you think most lawyers using AI at the firm perceive AI in a similar way to you?

How interested do you think they are in AI?

Do you think lawyers need to understand what AI is in order to use it effectively in their work?

Section 2 – Narrative Timeline

Turning now to your narrative timeline, talk me through the different influences you identified?

Possible probes:

What is your general attitude towards the use of AI at work? Do you have similar attitudes about technology more generally?

How did AI become part of your professional practice?

What goals is the use of AI designed to help realise? How is it meant to achieve this?

Section 3 – Reflective Tetrad

Turning now to the second part of the reflective exercise based on your experience of AI, talk me through the reflections you made in each of the four areas.

Possible probes:

How would you describe your overall experience of using AI at the firm?

Are you aware of any instances when the AI has failed to work as expected?

I'm interested to better understand your relationship with AI. I'm going to describe four common ways in which people can relate to technology. I'd be interested to know if one or more of them describes how you relate to AI, or whether you would describe your relationship differently.

Embodiment relations are when technology becomes an extension of our own human body. Examples might include using a pen to write or a car to drive.

Hermeneutic relations are when the technology must be interpreted or 'read' by us to generate meaning. Examples of this would be reading a thermometer or map.

Alterity relations are when the technology is experienced as independent from us, and its behaviour is hard to predict. An example could be interacting with a robot; but it can also

describe when we are learning to use a new technology; or when a technology breaks down.

Background relations are when the technology disappears from our conscious view but remains in contact with us. Examples include infrastructure such as heating systems and CCTV cameras.

Possible probes:

Does the type of relationship differ amongst the AI technologies you use?

Is your relationship with AI the same as with the other technologies you use as part of your work? For example, email or MS Word.

Section 4 – Trust & AI

Do you trust the AI technologies you use?

Possible probes:

If so, what is that trust based on?

If you don't trust it, why do you continue to use it?

Do you perceive AI to be more or less trustworthy than other technology?

How did your initial trust in AI develop?

Possible probes:

Prior to gaining first-hand experience of using the technology what factors led you to trust/distrust AI?

How has your trust in AI developed over time?

Possible probes:

How long did it take you to develop trust? What factors most influenced this?

Was this process of trust development similar to other technologies you have used?

Section 5 – AI & Professional identity

Our professional identity describes how we perceive ourselves within our occupational context and how we communicate this to others. Has AI had an impact on your professional identity [as a lawyer]?

Possible probes: How did this shift in identity take place?

Section 6 - Close

Is there anything I have not asked you about that you think is important about this topic?

Are there any questions you would like to ask me or questions you would like to revisit?

Are there any people in your organisation that you would recommend I approach to take part in the research project?

Appendix 2: Information sheet

Research Project Title: How lawyers experience the implementation of artificial intelligence in the workplace.

Invitation to participate

You are being invited to take part in a research project. Before you decide whether or not to participate, it is important for you to understand why the research is being done and what it will involve. Please take time to read the following information carefully and discuss it with others if you wish. Please ask if there is anything that is not clear or if you would like more information. Take time to decide whether or not you wish to take part. Thank you for reading this.

What is the project's purpose?

You are being asked to take part in a research project investigating lawyers' experience of working with artificial intelligence (AI) technologies, as part of the lead researcher's (Edward Walker) PhD studies. If you do not wish to take part in the study this will not affect his academic assessment, so do not feel obliged to volunteer for this reason or any other.

1. Why have I been chosen?

This project aims to reflect the varying experiences of different stakeholder groups who have personal experience of using AI technologies in their work. To ensure that the views of all relevant stakeholder groups are considered, you (and several of your colleagues) are being invited to participate in the research project. By leveraging the different kinds of knowledge held by diverse stakeholders, the findings of the research should be more useful than when scholars or practitioners work alone.

2. Do I have to take part?

It is up to you to decide whether or not to take part. If you do decide to take part, you will be given this information sheet to keep (and be asked to sign a consent form). You can still withdraw at any time prior to the publication of the research without any negative consequences. You do not have to give a reason. If you wish to withdraw from the research please contact Edward Walker.

3. What will happen to me if I take part? What do I have to do?

Data collection is anticipated to be completed by 31/10/2022. If you choose to take part in the research you will be invited to share your general perceptions of the use of AI in the legal profession, as well as your own specific, first-hand experiences of using AI at work.

Data collection will be conducted remotely to remain compliant with any COVID-related guidelines that are in place. Various methods of data collection will be used, including a reflective exercise and interviews. Participants will be asked to briefly record the factors that have influenced their attitude towards AI and how AI affects their own work. The amount of time dedicated to the reflective exercise is at the discretion of the participant, but it estimated to be no longer than 60 minutes in total. Interviews will be conducted online using video-conferencing software and will typically last 60-90 minutes. These will take place after the completion of the reflective exercise and will allow for a discussion of the reflective exercise, as well as a more general conversation about the experience of using AI.

The digital recordings of the data collected (both the reflective exercise and interview) will be electronically transcribed and used for analysis. As part of this process your words may be quoted in writing in publications, reports, web pages, and other research outputs on an anonymous basis, unless you request that your name is included. No other use will be made of the recordings without your written permission, and no one outside the project will be allowed access to the original non-anonymised recordings.

4. Will I be recorded, and how will the recorded media be used?

The recordings of the video interview and content of the reflective exercise from this research will be used only for analysis. No other use will be made of them without your written permission, and no one outside the project will be allowed access to the original recordings, other than for the process of transcription.

5. Will my taking part in this project be kept confidential?

All personal information collected during the interview and reflective exercise will be kept strictly confidential and will only be accessible to members of the research team.

Data will be pseudonymised prior to analysis, meaning it will not be possible to identify named individuals in any reports or publications unless they have given their explicit consent for this. Where future sharing of data with other researchers takes place (e.g. by making it available in a data archive), no personal details will be included, unless the participant explicitly requests this.

For clarity, the consequences to participants who elect not to remain confidential is that their name and job role (but no further personal details) will be attributable to any data they provide during the research. This would mean their views on the topic of AI would be publicly known.

6. What are the possible disadvantages and risks of taking part?

Given the topic of research and chosen methods for data collections, no foreseeable discomforts, disadvantages, and risks are anticipated from taking part in this research.

7. What are the possible benefits of taking part?

By participating in the research, you will gain the opportunity to share your experiences and perceptions of AI in the workplace. Through this you will be helping shape future research that has the potential to positively influence the experience of people working in professional service firms.

8. What is the legal basis for processing my personal data?

According to data protection legislation, we are required to inform you that the legal basis we are applying in order to process your personal data is that, 'processing is necessary for the performance of a task carried out in the public interest' (Article 6(1)(e)). Further information can be found in the University's Privacy Notice <https://www.sheffield.ac.uk/govern/data-protection/privacy/general>.

9. What will happen to the data collected, and the results of the research project?

Identifiable data will be available only to the lead researcher and their supervisors. It is anticipated that the preliminary results of the research will be made available to participants and other interested parties in 2022. The research findings, including anonymised extracts of any data collected by the researcher, may be quoted in publications, reports, web pages, and other research outputs.

Identifiable personal data (including the key which links an individual to the data they provided) will be destroyed as soon as possible, once the research and related publications have been completed. Due to the nature of this research it is very likely that other researchers will find the data collected to be useful in answering future research questions. Following completion of all phases of the research all anonymised data will be stored in ORDA, an online data repository managed by the University of Sheffield that can be accessed by other researchers. We will ask for your explicit consent for your data to be shared in this way.

10. Who is organising and funding the research?

The research is being funded by Sheffield University Management School.

11. Who is the Data Controller?

The University of Sheffield will act as the Data Controller for this study. This means that the University is responsible for looking after your information and using it properly.

12. Who has ethically reviewed the project?

This project has been ethically approved via the University of Sheffield's Ethics Review Procedure, as administered by the Management School (every academic department either administers the University's Ethics Review Procedure itself, internally within the department, or accesses the University's Ethics Review Procedure via a cognate, partner department). The University's Research Ethics Committee monitors the application and delivery of the University's Ethics Review Procedure across the University.

13. What if something goes wrong and I wish to complain about the research?

If you wish to make a complaint relating to the conduct of the lead researcher or the research project more, generally please contact the project supervisor. If the project supervisor is unable to respond to your complaint to your satisfaction, then you can contact the chair of the departmental ethics committee (please see below for further details).

If the complaint relates to the handling of personal data, information about how to raise a complaint can be found in the University's Privacy Notice: <https://www.sheffield.ac.uk/govern/data-protection/privacy/general>

14. Contact for further information

If you require further information or have any issues with the research do not hesitate to contact us for further information. You can contact the following people if you require further information or help:

Lead Researcher

Edward Walker

Email: [REDACTED]

Telephone: [REDACTED]

Project supervisor:

[REDACTED]
Email: [REDACTED]

Telephone: [REDACTED]

Chair of departmental research committee:

[REDACTED]

Email: [REDACTED]

Telephone: [REDACTED]

What happens next?

We hope this information has given you a better understanding of this research project. If you wish to take part in the research, please complete the consent form provided to you by the lead researcher and retain a copy of this Information Sheet for your future reference.

Thank you for considering participating.

Appendix 3: Consent form

<i>Please indicate your response in the appropriate boxes</i>	Yes	No
Taking Part in the Project		
I have read and understood the project information sheet dated 24.08.22 or the project has been fully explained to me. (If you answer No to this question please do not proceed with this consent form until you are fully aware of what your participation in the project will mean.)	<input checked="" type="checkbox"/>	<input type="checkbox"/>
I have been given the opportunity to ask questions about the project.	<input type="checkbox"/>	<input type="checkbox"/>
I agree to complete a reflective exercise of my experiences relating to the use of AI technology prior to being interviewed.	<input type="checkbox"/>	<input type="checkbox"/>
I agree to being video interviewed by the researcher on one or more occasions.	<input type="checkbox"/>	<input type="checkbox"/>
I agree to the interview being recorded, and a verbatim transcription made.	<input type="checkbox"/>	<input type="checkbox"/>
I understand that my taking part is voluntary and that I can withdraw from the research prior to the publication of its findings; I do not have to give any reasons for why I no longer want to take part and there will be no adverse consequences if I choose to withdraw.	<input type="checkbox"/>	<input type="checkbox"/>
How my information will be used during and after the project		
I understand my personal details such as name, phone number and email address etc. will not be revealed to people outside the project.	<input type="checkbox"/>	<input type="checkbox"/>
I understand that in any outputs of this research my identity will remain anonymous, unless I specifically request otherwise. This will be achieved by changing my name and disguising any details which may reveal my identity or the identity of people I speak about.	<input type="checkbox"/>	<input type="checkbox"/>
I understand and agree that disguised extracts of any data collected by the researcher from me may be quoted in publications, reports, web pages, and other research outputs.	<input type="checkbox"/>	<input type="checkbox"/>
I understand and agree that other authorised researchers will have access to this data only if they agree to preserve the confidentiality of the information as requested in this form.	<input type="checkbox"/>	<input type="checkbox"/>
I understand and agree that other authorised researchers may use my data in publications, reports, web pages, and other research outputs, only if they agree to preserve the confidentiality of the information as requested in this form.	<input type="checkbox"/>	<input type="checkbox"/>
I give permission for my anonymised data to be deposited in the ORDA depository so it can be used for future research and learning.	<input type="checkbox"/>	<input type="checkbox"/>
So that the information you provide can be used legally by the researchers		
I agree to assign the copyright I hold in any materials generated as part of this project to The University of Sheffield.	<input type="checkbox"/>	<input type="checkbox"/>

Name of participant [printed]

Signature

Date

Name of Researcher [printed]

Signature

Date

The template of this consent form has been approved by the University of Sheffield Research Ethics Committee and is available to view here: <https://www.sheffield.ac.uk/rs/ethicsandintegrity/ethicspolicy/further-guidance/homepage>

Appendix 4: Question list for organisational sponsors

These questions have been designed to help the research investigator (Edward Walker) gather information about the use of artificial intelligence at the firm. The information from this discussion will be used to help familiarise him with firm-specific practices and technology that participants may refer to in later interviews. Answers to these questions should be factual or reflect the dominant view held within the firm.

Q1. What technology is currently in use that the firm regards as being an example of artificial intelligence (AI)?

The following questions should be answered separately for each of the technologies identified above.

Q2. What in the firm's view makes this technology an example of AI?

Interviewer note: Probe in interview to determine whether any of the five assumptions linked to the tool perspective of technology are challenged by the technology.

Q3. What is the overall purpose of the AI? What wider service(s) is it embedded within?

Q4. How was the AI procured/developed?

Q5. Approximately how long has AI been in use at the firm?

Q6. Where is AI being used in the firm (i.e. which practice areas and locations)?

Q7. Who is using the software (i.e. what job roles)?

Q8. How widespread is the use of AI; how frequently does a typical user use AI-enabled services?

Q9. Is the use of AI mandatory or optional? What (if any) alternative options exist?

Q10. What support or training is given to individuals using AI?

Q11. How is AI's performance measured? How has it performed to date?

Q.12 Have there been any problems associated with the software?

Interviewer note: After discussing the above questions, request a demonstration of the AI technology that has been discussed.

Appendix 5: Participant instructions for reflective exercises

Pre-interview reflective exercise

Thank you for agreeing to take part in this research project. Prior to attending your interview with me, I would like you to complete a reflective exercise, details of which are provided below.

The exercise has been designed to help you identify the influences that have shaped your views about the use of artificial intelligence (AI) at work. The exercise involves two separate parts:

- Constructing a narrative timeline of key moments from the first time you became aware of AI being used in the workplace, to the present.
- Capturing your ideas about how AI affects your work in different ways.

It should initially take you about **30 minutes** to complete the above two exercises. You can also choose to update the exercise with additional data at any point prior to your interview.

When you are finished, please return the completed reflective exercise to [REDACTED] **at least one week** before your interview is scheduled to take place.

The reflective exercise will serve as a starting point for our interview and will allow for a richer discussion than would otherwise be possible.

If you have any questions relating to this reflective exercise, please do not hesitate to get in touch.

I look forward to our interview,

Edward

Edward Walker
Doctoral Researcher
Sheffield University Management School
[REDACTED]

www.sheffield.ac.uk/management

Part 1 – Narrative Timeline

Instructions for creating your narrative timeline

1. Use the blank worksheet provided below to create your narrative timeline.
2. Think about the different influences that have shaped your attitude and beliefs about the use of AI at work (rather than its use in other parts of your life). Place each influence on the timeline to reflect *when* they have influenced you; this could be a specific moment or an extended period. Don't worry if you can't remember precise dates, a rough estimate is fine.
3. Briefly *describe* each influence by capturing its salient points or related anecdotes. There will be an opportunity to discuss these in more detail in the interview.
4. Finally, please indicate on the timeline the point at which you first started using AI at work.

Advice when creating your narrative timeline

- If you have several different influences to record, you can construct your timeline across multiple pages if this makes it easier for you.
- Place positive influences above the timeline and negative influences below it.
- Record the intensity of the influence on you (e.g. high; medium; low)
- Pay particular attention to the period during which you have been using AI at work. Think about what has influenced you in relation to the specific AI software you use.
- Reflect on your continued use of AI between now and the interview, updating your timeline with your new insights.
- An example timeline for a junior lawyer new to AI is provided for guidance. Don't let this dictate how you construct your timeline, which is likely to be more detailed.

General categories of influences

You may find it useful to consider the following when creating your timeline:

Experiences – For example, how has your first-hand experience of using AI influenced you?

Objects – For example, how does the software design or its outputs influence you?

Individual agents – For example, does your line manager influence your attitude towards AI?

Groups – For example, how do the views of clients influence your attitude to AI?

Rules & practices – For example, do firm guidelines about the use of AI influence you?

Wider context – For example, how does wider society and the media influence you?

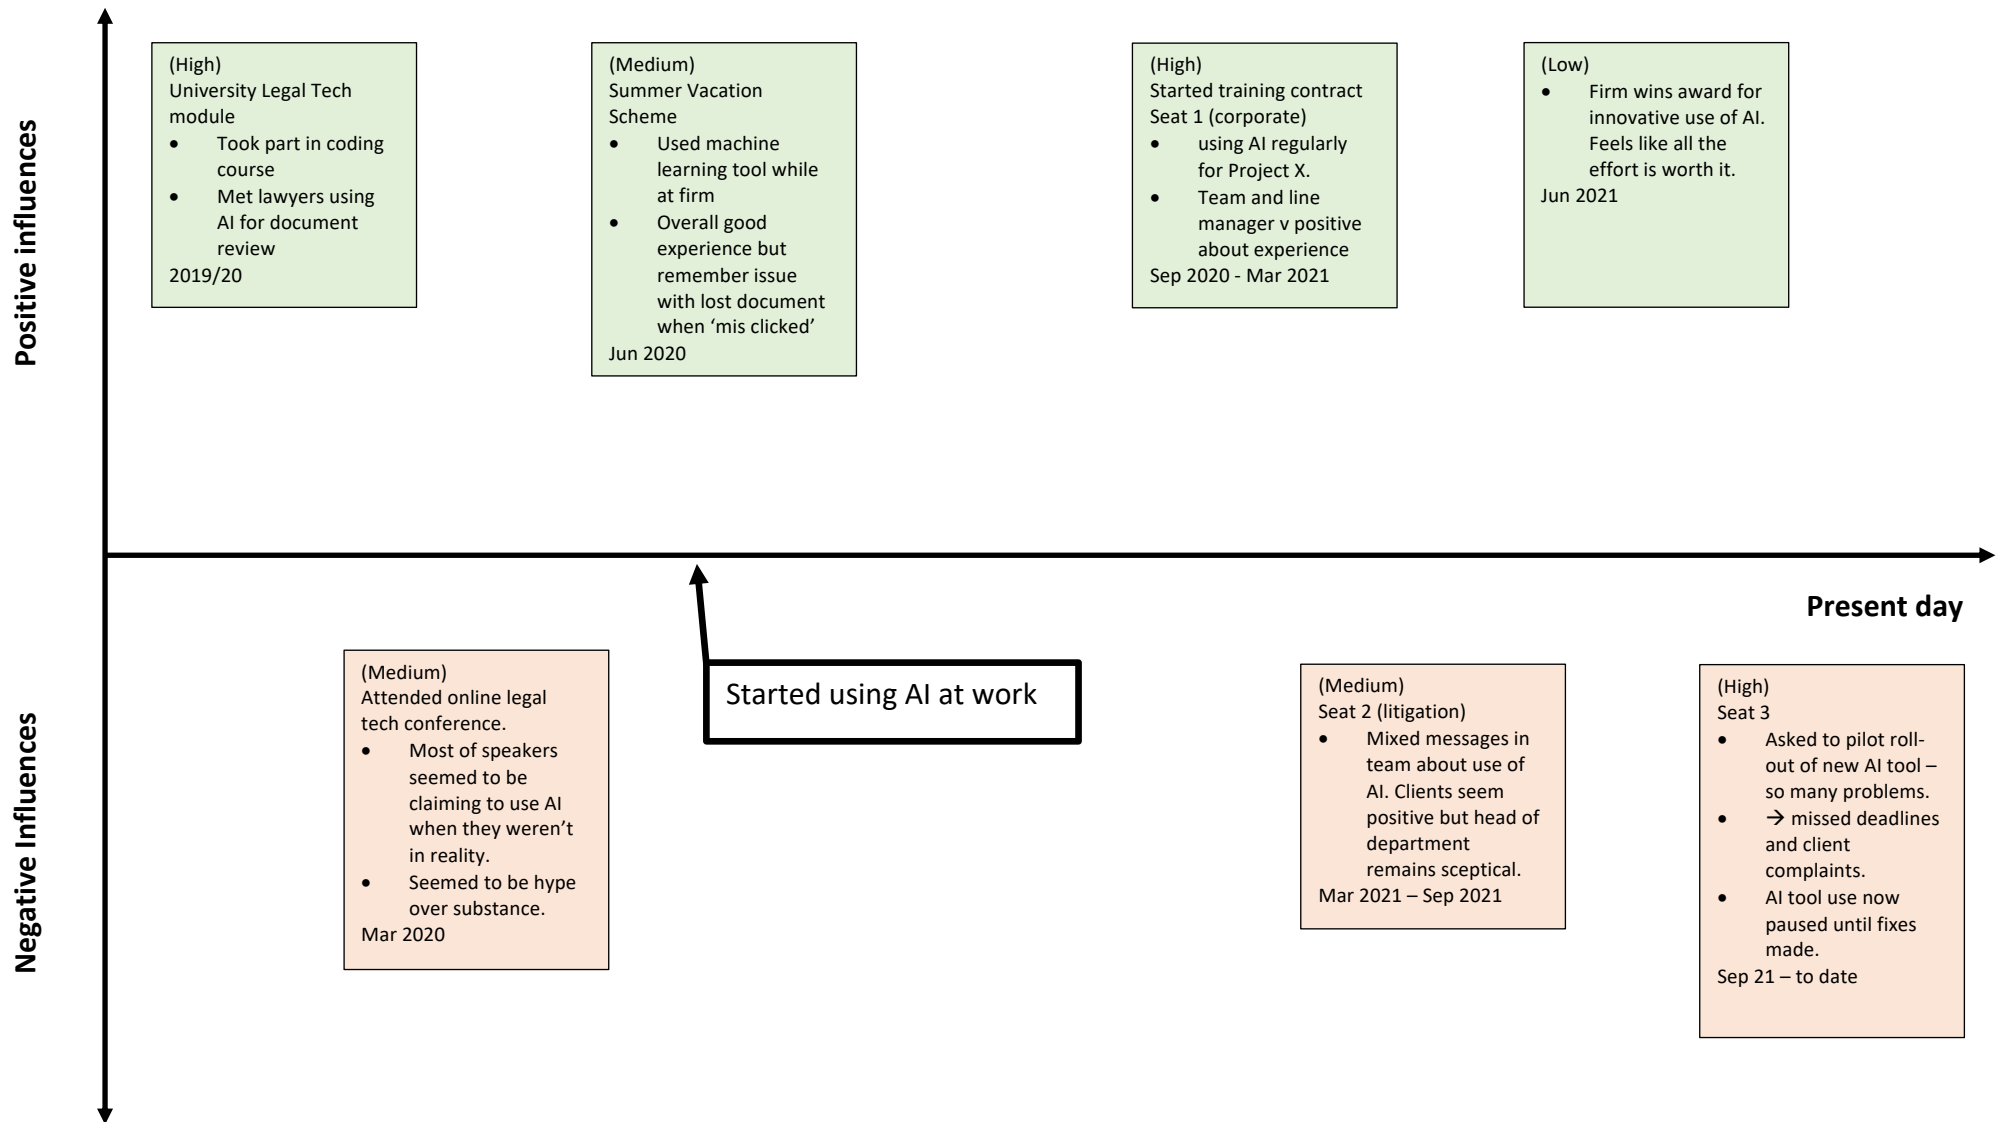
The list is not meant to be exhaustive, and there is no expectation that your timeline will contain examples from all six categories.

Useful definitions

An *attitude* is the way we think and feel about ourselves, another person, thing, or idea. Attitudes can be positive, negative, or neutral. With complex phenomena, such as AI, we may have varying attitudes towards different aspects of the phenomena. For example, we may have a positive attitude towards AI in general, but a negative attitude to a specific piece of AI software.

A *belief* is part of a system that includes our attitudes, plus our personal knowledge, experiences, opinions, morals, and other interpretive perceptions of the social world. For example, we may believe that combining the use of AI software with our personal judgement enhances the overall accuracy of our work.





Positive influences

(High)
 University Legal Tech module
 • took part in coding course
 • met lawyers using AI for document review
 |
 2019/2020

(Medium)
 Summer vacation scheme
 • used machine learning tool while at firm
 • overall good experience but remember issue with lost document when 'mis-clicked'
 |
 June 2020

(High)
 Started training contract
 Seat 1 - corporate using AI regularly for Project X. Team & line manager v positive about experience
 |
 September 2020 - March 2021

(Low)
 Firm wins award for innovative use of AI. Feels like all the effort is worth it.
 |
 June 2021

Negative Influences

(Medium) Attended online Legal Tech conference.
 • most of speakers seemed to be claiming to use AI when they weren't in reality - seemed to be hype over substance
 |
 March 2020

Started using AI at work

(Medium) March 2021 - Sep 2021
 Seat 2 - litigation mixed messages in team about use of AI clients seem positive but head of department remains skeptical.
 |

(High) |
 September 2021
 Seat 3 asked to pilot roll-out of new AI tool - so many problems -> missed deadlines & client complaints. Tool use now paused until fixes made.

Part 2 – How AI affects your work

Instructions for capturing your ideas

- Use the blank worksheet provided below to capture how AI affects your work in different ways.
- Based on your first-hand experience of using workplace AI, answer the following four questions about how your AI affects your work.
- An example of a completed worksheet has been provided for guidance. The example is based on how the use of GPS for navigation (a 'sat-nav') might affect an individual.

Questions

1. **Enhances** – What human capacities does your workplace AI enhance?
2. **Obsolesces** – What human capacities does your workplace AI diminish or render obsolete?
3. **Retrieves** – What previously obsolete tasks or capacities has your workplace AI made possible again?
4. **Reverses into** – What might happen if your workplace AI is used everywhere or to its maximum potential?

Enhances - What human capacities does your workplace AI enhance?

Obsolesces - What human capacities does your workplace AI diminish or render obsolete?

Retrieves - What previously obsolete tasks or capacities has your workplace AI made possible again?

Reverses into - What might happen if your workplace AI is used everywhere, or to its maximum potential?

Enhances - What human capacities does GPS enhance?

- My ability to travel from 'here' to anywhere.
- My confidence when travelling to new places.
- My ability to respond to unforeseen problems and minimise delays while travelling.

Obsolesces - What human capacities does GPS diminish or render obsolete?

- My ability to read a traditional map.
- My ability to imagine space in other ways e.g. bird's eye view.
- My ability to offer directions to visitors to my hometown.

Retrieves - What previously obsolete tasks or capacities has GPS made possible again?

- The ability to go and explore anywhere without the need for specialist maps and equipment, as people did in the past with just a compass.

Reverses into - What might happen if GPS is used everywhere, or to its maximum potential?

- GPS becomes the only means through which humans can determine where 'here' is.
- Humans lose the ability to navigate unaided and become lost if their GPS breaks.
- Humans start to doubt their own 'sense of direction' when making journeys.

Appendix 6: Experiential theme summary

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	Frequency	
Interaction with AI-enabled legal services	PR	LIM	CS	CS	LIM	PR	CS	LIM	CS	CS	CS	PR	PR	CS	CS	PR	CS	LIM	LIM	LIM			
Most legal professionals in this study feel their role is to provide interpretation and meaning within an AI-enabled legal service.	TRUE		TRUE	TRUE	TRUE		TRUE		TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	18
Most legal professionals in this study perceive AI as just one part of a wider process.			TRUE	TRUE	TRUE			TRUE		TRUE	TRUE		TRUE	TRUE				TRUE	TRUE	TRUE			11
Most legal professionals in this study do not feel more negatively about their role from using AI-enabled legal services.	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	21
Most legal professionals in this study have difficulty defining AI.		TRUE	TRUE					TRUE	TRUE	TRUE	TRUE			TRUE	TRUE				TRUE	TRUE	TRUE		11
Most legal professionals in this study believe the demands of clients drive the use of AI-enabled legal services.	TRUE	TRUE		TRUE	TRUE	TRUE	TRUE	TRUE	TRUE		TRUE				TRUE			TRUE		TRUE			12
Many legal professionals in this study believe competitor behaviour can drive the use of AI-enabled legal services.	TRUE				TRUE	TRUE	TRUE	TRUE		TRUE		TRUE		TRUE								TRUE	9
Some legal professionals in this study believe their firm promotes the use of AI-enabled legal services.		TRUE				TRUE				TRUE	TRUE	TRUE			TRUE			TRUE	TRUE				8
Most legal professionals in this study believe that AI-enabled legal services are necessary for certain kinds of legal task.	TRUE				TRUE	TRUE			TRUE	TRUE	TRUE	TRUE	TRUE			TRUE			TRUE		TRUE		11
Most producers of AI-enabled legal services believe them to be understandable.	TRUE					TRUE		TRUE					TRUE				TRUE		TRUE				6
Most consumers of AI-enabled legal services believe them to lack transparency.							TRUE	TRUE	TRUE	TRUE	TRUE			TRUE		TRUE			TRUE	TRUE	TRUE		10
Most producers report initial trust in AI-enabled legal services stems from the technology itself or someone with expert knowledge of it.					TRUE	TRUE						TRUE	TRUE				TRUE						5
Most consumers report initial trust in AI-enabled legal services stems from colleagues.			TRUE	TRUE					TRUE					TRUE				TRUE		TRUE	TRUE		7