

**ON THE BEHAVIOURAL PROFILING OF GAMBLERS
USING CRYPTOCURRENCY TRANSACTION DATA**

by Oliver J. Scholten

Doctor of Philosophy

University of York

Computer Science

July 2022

Abstract

Blockchain technologies enable a number of new ways to gamble online. Very little is known about engagement with one such new way of gambling: decentralised gambling applications, which provide simple casino games like dice rolls and coin flips. This is important as understanding engagement with any type of gambling is a crucial first step to assessing the risk of experiencing gambling related harm within the population. This thesis first surveys existing literature for methods of describing engagement in gambling, and then applies these methods to actual transaction data gathered from several decentralised gambling applications. This replication-oriented approach means results can be grounded against existing findings, and the descriptions of player engagement in this new domain have some context for comparison. It also means that descriptions can be tentatively mapped to similar scenarios, such as risk of experiencing gambling related harm in other studies. The results of several replication oriented studies presented herein find that engagement in the decentralised gambling domain is typically less than in comparable online casino games, but that a heavily involved subgroup is more involved. It also finds that engagement with gambling-like mechanisms in blockchain games is much less than in decentralised gambling applications, guiding future studies in gambling research away from blockchain games despite their mechanical similarities. Finally, behavioural groups in the decentralised gambling domain do not appear to be comparable with existing research in the centralised online casino game domain. The results of these studies provide a first look at engagement in this emerging domain, a comparative description with similar forms of gambling, and a description of behavioural groups, which provides essential context for further research to assess the scale of the risk of experiencing gambling related harm.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Research Questions	2
1.3	Thesis Outline	2
1.4	Scope	4
1.5	Contributions	5
2	Cryptocurrency Fundamentals	7
2.1	Cryptographic Primitives	7
2.1.1	Hash Functions	8
2.1.2	Asymmetric Cryptosystems	8
2.1.3	Digital Signatures	9
2.2	Blockchains	11
2.2.1	Data Structure	11
2.2.2	Distributed Computing	11
2.2.3	Consensus Mechanisms	12
2.3	Cryptocurrencies	13
2.3.1	Addresses & Wallets	14
2.3.2	Smart Contracts	17
2.3.3	Non-Fungible Tokens	19
2.3.4	Price Volatility	21
2.4	Decentralised Applications	23
2.4.1	Gambling	24
2.4.2	Gaming	25
2.5	Summary	26
3	Decoding Cryptocurrency Transactions	27
3.1	Accessing Cryptocurrency Transactions	28
3.1.1	Node Synchronisation	28
3.1.2	Public API's	29
3.2	Transaction Types	30
3.2.1	Normal Transactions	31
3.2.2	Internal Transactions	31

3.3	Identifying Function Calls	32
3.3.1	Solidity Programming	32
3.3.2	Gas Fees	33
3.3.3	Method Signatures	34
3.3.4	Parameter Counts & Values	36
3.4	Emulating Application Mechanics	37
3.5	Matching Payouts	38
3.5.1	Naive Chronological Stitch	39
3.5.2	Breadcrumb Trace	39
3.6	Summary	40
4	Player Behaviour Tracking	42
4.1	Types of Gambling	43
4.1.1	Mercantile versus Social Gambling	44
4.2	Types of Game	45
4.2.1	Games of Skill	45
4.2.2	Games of Chance	46
4.2.3	The Skill-Chance Spectrum	47
4.3	Data Types	47
4.3.1	Account Level	47
4.3.2	Transaction Level	48
4.4	Study 1: Systematic Review	49
4.4.1	Introduction	50
4.4.2	Method	51
4.4.3	Results	54
4.4.4	Discussion	57
4.5	Summary	62
5	Understanding Behavioural Measures	64
5.1	Existing Research	65
5.1.1	Study Variation	66
5.2	Behavioural Domains	68
5.2.1	Time Domain	68
5.2.2	Bet Domain	78
5.2.3	Loss Domain	83
5.2.4	Risk Domain	86
5.2.5	Miscellaneous Domain	87
5.3	Behavioural Measure Clustering	88
5.4	Cost of Gambling	90
5.5	Summary	91

6	Decentralised Gambling Application Prevalence	93
6.1	Study 2: Gambling Application Prevalence	94
6.1.1	Introduction	94
6.1.2	Method	95
6.1.3	Results	97
6.1.4	Discussion	102
6.1.5	Conclusion	106
6.2	Study 3: Blockchain Game Prevalence	106
6.2.1	Introduction	106
6.2.2	Method	107
6.2.3	Results	110
6.2.4	Discussion	111
6.2.5	Conclusion	116
6.3	Summary	117
7	Inside Decentralised Gambling Applications	118
7.1	Introduction	119
7.2	Application Architectures	119
7.2.1	Case Study: dice2.win	120
7.2.2	Case Study: Etheroll.com	123
7.2.3	Case Study: FCK.com	127
7.3	Study 4: Behavioural Distributions in Gambling Applications	129
7.3.1	Introduction	129
7.3.2	Method	130
7.3.3	Results	134
7.3.4	Discussion	144
7.3.5	Conclusion	149
7.4	Summary	150
8	Inside Blockchain Games	152
8.1	Introduction	153
8.2	Application Architectures	154
8.2.1	Randomised Reward Function Calls	154
8.2.2	Case Study: CryptoKitties	155
8.3	Study 5: Behavioural Distributions in Blockchain Games	157
8.3.1	Introduction	157
8.3.2	Method	159
8.3.3	Results	161
8.3.4	Discussion	165
8.3.5	Conclusion	168
8.4	Summary	169

9 Identifying Behavioural Groups	170
9.1 Study 6: Behavioural Groups in Decentralised Gambling . . .	171
9.1.1 Introduction	171
9.1.2 Method	174
9.1.3 Results	177
9.1.4 Discussion	181
9.2 Summary	185
10 Conclusion	186
10.1 Contributions	187
10.1.1 Behavioural Measures	187
10.1.2 Decentralised Application Prevalence	187
10.1.3 Decentralised Gambling Behaviours	188
10.1.4 Blockchain Gaming Behaviours	189
10.1.5 Behavioural Groups in Decentralised Gambling Applications	189
10.2 Limitations	190
10.3 Further Work	192
A Study 2 Raw Plots	194
B Study 4 Additional Results	197
Bibliography	206

List of Tables

4.1	Search terms used in Chagas and Gomes’ systematic review [1] (top) and the search terms used in the systematic review for this thesis (bottom). The key changes are the inclusion of the British spelling of ‘behavioural’, and ‘machine learning’ over ‘player card’ and ‘loyalty card’ to capture more technically oriented papers.	51
7.1	A sample of dice2win smart contract ABI methods and their functionality (verified from address: <code>0xD1CEeeeeee83F8bCF3BEDad437202b6154E9F5405</code>).122	
7.2	A sample of Etheroll smart contract ABI methods (verified from address: <code>0xf478c8Bc5448236d52067c96F8f4C8376E62Fa8f</code>). Method in bold represents the core player action of interest used in further studies.	125
7.3	A sample of FCK smart contract ABI methods (verified from address: <code>0x9999999C60566e0a78DF17F71886333E1dACE0BAE</code>).128	
7.4	Smart contract addresses for each decentralised gambling application used in this study.	131
7.5	Meta data for each application gathered as part of this study. Bet and Payout values are given in ETH, and starting and ending blocks and dates represent the time window from which transactions were gathered.	131
7.6	Two sample Kolmogorov-Smirnov test results for player durations across all provider-game combinations. † denotes a significant result ($p < 0.01$) and coefficients greater than 0.35 are highlighted. Key: d2w = dice2.win, fck = FCK.com, eroll = Etheroll.com, cf = coin flip, sd = single dice roll, dd = double dice roll, oh = 1–100 roll.	137

7.7	Gambling behaviour of 10,357 decentralised gambling application players including a one sample Kolmogorov-Smirnov (K-S) test for normality. All K-S test statistic values are significant at the $p < 0.01$ level, STD = standard deviation, IQR = inter-quartile range.	138
7.8	Non-parametric Spearman rank-order correlations between all behavioural measures for decentralised gambling application players. All values are significant at the $p < \mathbf{0.01}$ level. Coefficients of magnitude greater than 0.70 are highlighted. .	139
7.9	Non-parametric descriptive statistics of the behavioural measures for the top 5% most heavily involved bettors by total amount wagered, and the other 95% of players. All one sample K-S test statistic values are significant at the $p < 0.01$ level indicating the data for each measure is non-normally distributed.	141
7.10	Gambling behaviour of extreme 5 and 95% subgroups of casino bettors, reprinted from LaBrie et al's 2008 study [2].	142
8.1	Behavioural measures applied to CryptoKitties' <code>breedWithAuto</code> function calls.	162
8.2	Spearman's correlation coefficient scores between each pair of the six measures computed across CryptoKitties <code>breedWithAuto</code> function calls. Key: dur = duration, freq = frequency percent, num = number of bets, mday = mean number of bets per day, msize = mean bet size, tot = total amount wagered. All values are significant at $p < 0.05$	163
8.3	Descriptive statistics on the set of behavioural measures computed (n=21,897), split by heavy involvement (top 5%) by total amount wagered. Key: dur = duration, freq = frequency percent, num = number of bets, mday = mean number of bets per day, msize = mean bet size, tot = total amount wagered.	164
8.4	Gambling behaviour of extreme 5 and 95% subgroups of casino bettors, reprinted from LaBrie et al's 2008 study [2].	165
9.1	Behavioural measures which are used as input to the clustering algorithm (as used by Braverman and Shaffer [3]), computed using players' first month of transaction data.	175

9.2	Behavioural measures which are used to describe player's behaviours, computed on both the first month of data and entire careers of data - these are not used as input to the clustering algorithm.	176
9.3	Standardised behavioural measure (number of standard deviations from mean) scores of cluster centroids using a five cluster solution across the first month of cryptocurrency gambling activity (n=3,870). Extreme values are displayed in bold.	178
9.4	Clustering results reproduced from the original study for reference [3]. Extreme values are displayed in bold.	178
9.5	Behavioural measures describing the first month of gambling activity for each cluster identified using the k-means algorithm. Notable values are displayed in bold.	179
9.6	Behavioural measures describing the entire careers of gambling activity for players in each cluster identified by the k-means algorithm. Notable values are displayed in bold.	181
B.1	Behavioural measure scores across the entire sample.	197
B.2	Spearman's rank correlation coefficient between behavioural measures across all players.	198
B.3	Heavily involved and non-heavily involved split across all players. Total player counts are lower than the study presented in Chapter 7 as although the data covers a longer period of time, players with duration less than 4 (9,241 out of 13,569) have been removed.	198

List of Figures

1.1	The structure of the chapters in this thesis, with arrows indicating logical dependencies between chapters.	3
2.1	Visual example of how the output of a hash function on a piece of data, and the same data which has been maliciously modified (red boxes signify malicious changes), are different.	9
2.2	Visualisation of how the three functions of a digital signature scheme (key generation, signing function, and verification function) interact to provide authenticity to a given piece of data using a public-private key pair.	10
2.3	A simple blockchain with a genesis block and three appended blocks. Each appended block contains a hash of the previous block, forming a singly linked append-only chain.	11
2.4	An example NFT from the Bored Ape Yacht Club collection (See https://boredapeyachtclub.com/ , accessed 01/02/2022). This particular ape (See https://opensea.io/assets/0xbc4ca0eda7647a8ab7c2061c2e1188585 for details, accessed 01/02/2022) sold on October 19 2021 for 696.969ETH, which at the time of writing is equal to approximately £1.4m (time of sale approximately £1.9m).	20
2.5	The price of ETH (the native currency of the Ethereum network) over time, see https://www.coingecko.com/en/coins/ethereum , accessed 04/11/2021.	22
2.6	Uniswap user interface example converting from one cryptocurrency (ETH) into another (MANA).	23

3.1	A simplified representation of the difference between normal and internal transactions, with a human user and two smart contracts. Both the human user and smart contracts have cryptocurrency addresses.	30
3.2	Example of a basic function whose source code (in green) is compiled and encoded (in yellow) when stored on the blockchain. It is clear that in cases where the source code for a given contract is unavailable, determining its functions and capabilities becomes a reverse engineering problem.	32
3.3	The structure of an Ethereum transaction's input field on a method with one parameter.	35
3.4	Example data set of transaction input fields showing that the function being called (e.g. <code>place_bet()</code>) can be identified by the signature at the start of the input field.	35
4.1	Each of the stages in the systematic review method, taking input collections of papers from Chagas & Gomes review, plus the up to date Web Of Science output. The result of this pipeline is a globally unique collection of behavioural measures used in gambling research.	53
4.2	Globally unique collection of behavioural measures identified in Study I in this thesis. Colour coding indicates the domain in which each behavioural measure resides, this is based on the information needed to compute each measure - a concept expanded upon below in Section 5.1. Greyed out measures required contextual information (time zone), and arrows signify a computational dependency between measures, e.g. Total number of sessions is required to compute amount wagered per session.	58
4.3	The distribution of behavioural measures used in studies uncovered by the systematic review. It is important to note that the raw number of times a behavioural measure is used in the literature is not necessarily and indicator of its validity, but acts as a useful starting point for further investigation.	59

4.4	Each of the unique behavioural measures uncovered by the systematic review, plotted against the studies they were used in. Note that in some instances in text a measure may be referenced as being used in two studies but appears as a three in this figure. This is due to variants or transformations of the same measure being used in a single study.	60
6.1	Usage metrics across all 147 decentralised gambling applications gathered as part of this thesis.	99
6.2	Usage metrics of the top 20 decentralised gambling applications in the data sample gathered for this study.	101
6.3	Usage metrics across all 260 blockchain games applications gathered as part of this thesis.	110
6.4	Usage metrics of the top 20 blockchain games in the data sample gathered for this study.	112
7.1	The dice2win application's bitwise roll under operation, taken from source code verified by Etherscan at address 0xD1CEeeee83F8bCF3BEDad437202b6154E	
7.2	dice2win application user interfaces available at https://dice2.win , accessed 11/08/2021.	124
7.3	Etheroll application user interface, available at https://etheroll.com accessed 13/09/2021.	126
7.4	Source code of the bitwise roll under operation found in the FCK smart contract. Source code verified by Etherscan from address 0x999999C60566e0a78DF17F71886333E1dACE0BAE.	128
7.5	Provider-game combinations, including unique address and bet counts taken forward to the final player transaction set.	130
7.6	Transaction data gathering timelines for each of the three decentralised gambling applications studied. FCK ceased operation in July 2019.	132
7.7	Cumulative value of bets placed through each application individually, and all applications combined over the period studied. The data and code used to create this figure is available at https://osf.io/8bfyj/	132

7.8	Theoretical differences in distributions of behavioural measures between human and non-human players. The second spike is created when multiple addresses transact in the same way, e.g. using by using computer code.	135
8.1	The general form of randomised reward mechanisms described by Nielsen and Grabarczyk [4].	154
8.2	Description of CryptoKitties, taken from https://www.cryptokitties.co/ , accessed 06/07/2021.	155
8.3	Distribution of the behavioural measure of duration within the blockchain games player sample (n=21,897). Note that the x-axis is discontinuous to highlight the extreme nature of the distribution.	162
9.1	Total value transacted (ETH) by all bets placed in the data sample (n=14,462, tx=2,336,468).	175
9.2	Mean inertia scores across repeated clustering solutions in the data sample (first month of gambling, repeats per k=30, n=3,870).	177
A.1	Usage metrics across all 147 decentralised gambling applications gathered as part of this thesis.	195
A.2	Usage metrics of the top 10 decentralised gambling applications in the data sample gathered for this study.	196

Acknowledgements

This thesis could only exist with the support of my supervisors James Walker and David Zendle. Thank you for the many discussions, suggestions, and debates over the last few years which have massively improved this work. Many thanks also to Ken Hawick, whose mentorship and belief in me as an undergraduate computer scientist gave me the direction and confidence to pursue postgraduate research.

Thank you to my colleagues and friends in the IGGI doctoral training centre for all of the office discussions, socials, and support over the years. Ryan in particular for being there from the early days at the University of Hull and following me everywhere I went from then on.

Thank you to my mum and sister for supporting me, and for helping me see the world outside of this research.

Finally, to any aspiring computer scientists reading this, I want you to know that even the most challenging topics, from cryptography and reverse engineering, to protocol analysis and high performance computing, can be mastered with enough tenacity, determination, and patience.

Good luck.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References. Some of the material in this thesis has previously been published in the following papers:

Scholten, Hughes, Deterding, Drachen, Walker, and Zendle - *Ethereum crypto-games: Mechanics, prevalence, and gambling similarities*, in: Proceedings of the Annual Symposium on Computer-Human Interaction in Play 2019. ACM, pp. 379-389

Scholten, Zendle, and Walker - *Inside the decentralised casino: A longitudinal study of actual cryptocurrency gambling transactions*, in: Public Library of Science (PLoS One) 2020, Volume 15 Issue 10.

Some material was also presented at the following conference:

Scholten - *Decentralised Gambling: Data Abundance and Technical Concerns*, in: Current Advances in Gambling Research Conference 2019.

Chapter 1

Introduction

*“Computer science is no more about computers
than astronomy is about telescopes.”*

Edsger Dijkstra

Cryptocurrency technologies enable a number of new ways to gamble online [5]. One such way is through the use of decentralised applications, which process payments and calculate game outcomes using code executed within a cryptocurrency blockchain. Recent advances in cryptocurrency technology mean that several such applications now exist, and provide simple casino games such as dice rolls and coin flips to a global audience. The technical properties of such applications means that they are resistant to regulatory actions commonly taken against unlicensed operators [6]. Additionally, their relative youth, in combination with their technical sophistication, mean that academic research into such applications has been limited [7]. This means that the way in which players engage with these applications is largely unknown, despite its importance to policy decisions and consumer protection efforts. This is the problem that this thesis aims to address.

1.1 Motivation

A small portion of players who engage in gambling experience some form of gambling related harm [8]. By understanding how people engage with different types of gambling, and how these types of engagement relate to experiencing such harms, existing studies have provided essential evidence to policy makers and regulators with the hope of mitigating these potential

harms. A foundational step in performing such studies is understanding how players engage with different types of gambling at the population level. Developing an understanding of player engagement in cryptocurrency gambling through decentralised applications is the first step on the long journey of reducing harms in this emerging domain.

1.2 Research Questions

This thesis aims to establish an understanding of player behaviours in the cryptocurrency gambling domain. Specifically, it focuses on the use of decentralised applications which offer simple casino games like dice rolls and coin flips. This aim can be broken down into the following set of research questions;

1. Which cryptocurrency transaction data be used for gambling research?
 - (a) How can this data be accessed?
 - (b) Does this data need preprocessing?
2. Which analytical methods can be meaningfully applied?
 - (a) Which behavioural measures can be computed?
 - (b) How can they be used?
3. How prevalent are decentralised gambling applications?
 - (a) Which applications may provide the most data?
4. What are player behaviours in these applications?
 - (a) How do they compare to other forms of gambling?
5. Which behavioural groups exist in this domain?
 - (a) How prevalent are they within the population?
 - (b) How do they compare to other forms of gambling?

1.3 Thesis Outline

The research presented in this thesis covers six studies which broadly map to the research questions outlined above. These are;

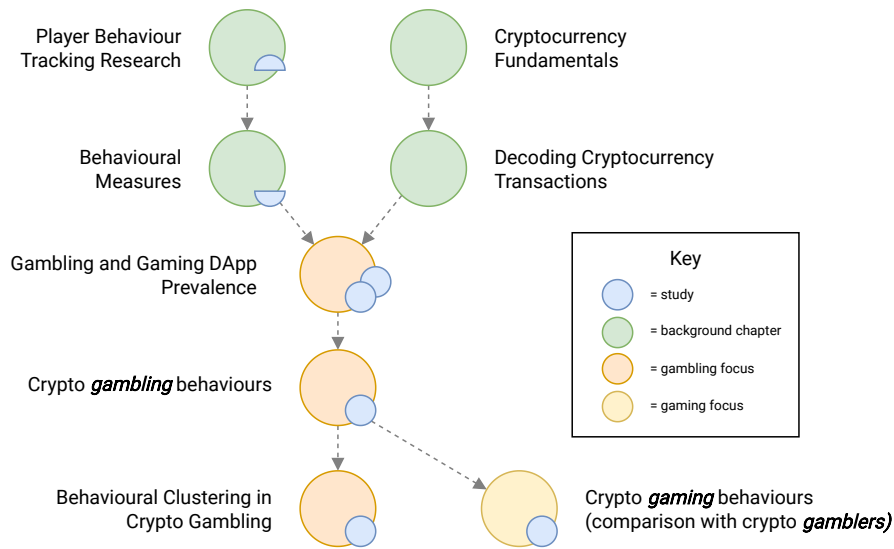


Figure 1.1: The structure of the chapters in this thesis, with arrows indicating logical dependencies between chapters.

1. A Systematic Review of Behavioural Measures used in Player Behaviour Tracking Research
2. Decentralised Gambling Application Prevalence
3. Blockchain Game Prevalence
4. Behavioural Distributions in Decentralised Gambling Applications
5. Behavioural Distributions in Blockchain Games
6. Behavioural Groups in Decentralised Gambling Applications

These studies are distributed throughout this thesis, following the structure presented in Figure 1.1. From this list of studies and structure, it is clear that a portion of this thesis is dedicated to blockchain games, rather than decentralised gambling applications. This is because such games can contain randomised reward mechanisms, which are mechanically similar to the simple casino games offered by the gambling applications [9]. The exact similarities are expanded upon in Chapters 6 and 8, but in the context of this thesis such mechanisms are considered a unique form of cryptocurrency gambling via a special case of decentralised gambling applications, so are therefore of interest.

1.4 Scope

This thesis does not explore all types of gambling using cryptocurrency [5]. Instead, it focuses exclusively on transactions to and from decentralised gambling applications, which execute bet placement functions atop a cryptocurrency network (see Chapter 2). This means that although the work in this thesis develops our understanding of the use of these applications in particular, the findings may not apply more broadly to other forms of cryptocurrency gambling such as gambling using cryptocurrency as stake in a regular casino. Additionally, this thesis focuses on simple casino games such as dice rolls and coin flips, as these are the most prevalent in available applications. This means that findings may not generalise to other types of gambling activity within cryptocurrency gambling.

Each of the studies in this thesis are also limited to the data available on the Ethereum cryptocurrency blockchain, chosen as it is the largest (by market capitalisation) and oldest network which explicitly enables the creation of decentralised gambling applications. While other cryptocurrencies do exist which contain the functionality needed to create such applications, these other currencies are not considered in scope of this thesis. This limitation is a result of the finite hours available to study different cryptocurrencies and applications, and means that while this work presents the beginning of academic research into such applications, much more can be done in the future.

In addition to this thesis' focus being on the Ethereum network, all subsequent analyses retain the native currency of the Ethereum network (ETH) for reporting. This means that throughout this thesis, the values of bets, payouts, and any derivative values are all described in terms of ETH rather than USD, GBP, or any other fiat currency. This approach has been taken over any 'real value' conversions (to USD for example) as in order to make such a conversion, the fundamental assumption that gamblers are buying, gambling, and then cashing out, their cryptocurrency close to the point of gambling must be met, due to the inherent volatility of cryptocurrencies in general and ETH in particular. Given the limited research into the use of decentralised gambling applications, the exact timings regarding player purchasing and selling cryptocurrency remains unknown. Until a description of these behaviours specific to the applications explored in

this thesis are published, it is difficult to commit to such a hard assumption and therefore more appropriate to retain the ‘native’ currency of ETH.

1.5 Contributions

The novel contributions presented in this thesis include parts of work previously published in peer reviewed conferences and an academic journal. Study 1 focuses on charting the development of the field of player behaviour tracking. This study reveals that despite consistent innovation across the field in terms of creating new behavioural measures, a core set of established measures exist which have been broadly applied across different gambling domains. A second key finding of this study is that in order to understand gambling in new domains, this set of established measures can be applied using transaction data alone in order to generate a meaningful description of behaviours upon which further work can be built.

Studies 2 and 3 focus on research question 3, and present both an overview of the scale of the usage of decentralised applications, and candidate applications whose available data invites analysis in a resource-efficient way. These studies each provide broader context for the later studies in this thesis, and provide empirical support for the choice of applications in these studies.

Study 4 applies the findings from Study 1 to the data gathered from leading decentralised gambling applications uncovered in Study 2. This presents the first ever behavioural description of players in this domain, and suggests that while the majority of players in this emerging domain do not engage with these applications very heavily, a small portion of players do spend extreme amounts. This is compared against similar studies in different domains, building a picture of the differences and similarities in behaviours between these domains.

Study 5 applies the findings from Study 1 to specific gambling-like mechanisms found in blockchain games. This reveals that the way in which players engage with these gambling-like mechanisms is very different from how players engage with comparable simple casino games, specifically that their financial involvement is far lower. These results suggest, at least in the cryptocurrency domain, that despite growing evidence of a link between problem gambling and engagement with such mechanisms, more research is needed to understand how/if this link actually affects player behaviours.

Study 6 extends Study 4 by applying behavioural clustering techniques in order to understand the distribution of behaviours within the population. This study finds that although a single ‘high activity, high variability’ group is not empirically identified as in comparable research [3], several groups exhibit extreme behaviours which suggest further research is required.

These contributions are important for the field of player behaviour tracking, as they demonstrate that cryptocurrency transaction data is in-fact useful for academic research, and suggest that the use of decentralised gambling applications is a unique method of gambling worth academic inquiry in its own right. As the use of cryptocurrency technology for gambling matures, one can expect to see more advanced casino games, sports betting, and other gambling activities provided in this way. This means that although this thesis focuses on simple casino games, its methodology may be later applied to understand and compare behaviours across other cryptocurrency gambling activities. A more detailed discussion of these contributions with reference to the research questions posed in Section 1.2 is deferred to Chapter 10.

Chapter 2

Cryptocurrency Fundamentals

“Cryptography is the ultimate form of non-violent direct action.”

Julian Assange

Cryptocurrencies are an advanced technology which make use of a number of different cryptographic and computational concepts. This chapter introduces these concepts, providing a necessary technical background to the subsequent chapters in this thesis. Section 2.1 introduces the necessary cryptographic primitives for understanding subsequent sections in this chapter. Section 2.2 then introduces blockchains - a type of data structure used in all cryptocurrencies. Section 2.2.2 then introduces distributed computing, with subsections describing how this concept can be combined with blockchains to provide distributed storage, and how this storage can be synchronised across the network. Section 2.3 then describes how cryptographic primitives, blockchain data structures, and distributed computing can be combined to create cryptocurrencies. Finally, Section 2.4 describes how applications can be built atop these cryptocurrencies, followed by a summary.

2.1 Cryptographic Primitives

Cryptographic primitives are a basic building block for higher-level cryptographic algorithms [10]. A basic knowledge of several such primitives is

necessary in order to understand how they are used to achieve different design goals in both blockchains and cryptocurrencies. The primitives important to discussion of cryptocurrencies are hash functions, asymmetric cryptosystems, and digital signatures.

2.1.1 Hash Functions

A hash function is an algorithm which takes an arbitrary length string and compresses it into a fixed length output [11]. These functions are also ‘one-way’, meaning that the data used as input cannot be determined from the output (a property known as preimage resistance). Further, for many hash functions a desired input cannot be computed which has the same output as a given output (known as second preimage resistance). Examples of hash functions include SHA-256 [12], Keccak-256 [13], and MD5 [14]. These functions are deterministic, which means that for any given input the output is always the same. This is a desirable property because for a given piece of data they can be used to generate a fixed length unique summary or fingerprint of that data. This can be used to provide the security notion of integrity, which means that if given a purported piece of data, and its original hash, one can cryptographically confirm that the data in question has not been maliciously modified¹. Figure 2.1 shows how a hash function can be used to compare two ‘copies’ of a piece data; an original (left), and one which has been maliciously modified (right). The ability of a hash function to provide a fixed length, unique summary of a piece of data, which can be used to check its integrity, makes them very useful in data structures where integrity is required (See Section 2.2).

2.1.2 Asymmetric Cryptosystems

Asymmetric cryptosystems, or public-key cryptosystems, are systems which use a pair of cryptographic keys to perform cryptographic tasks like encryption and digital signing [15]. The central component of these systems is a key generation algorithm, which for a given input (typically a large random number) generates two mathematically related numbers, called keys. These keys are related in that data encrypted using one key (public key) can be

¹Assuming a secure channel is available to share the hash, and except in the case of a collision, which in practical terms is extremely unlikely




Figure 2.1: Visual example of how the output of a hash function on a piece of data, and the same data which has been maliciously modified (red boxes signify malicious changes), are different.

decrypted by the other (private key). When used for encryption, asymmetric cryptosystems can be used to provide confidentiality - that only the intended recipient can decipher a message. Asymmetric cryptosystems are most relevant for their use in creating digital signatures, which combine the unique summary capabilities of hashing algorithms above with asymmetric encryption to create a way to ensure that not only has a message not been maliciously modified, but that it is indeed from the original sender.

2.1.3 Digital Signatures

Digital signatures are a way to create unique summaries of data similar to hash functions, which can then be verified to have been created by a specific person [16]. They achieve this by implementing a key generation algorithm, a signing algorithm, and a verification algorithm, the usage for which is shown in Figure 2.2. Examples of digital signature algorithms are RSA [17], ElGamal [18], and NTRUSign [19].

The ability of digital signatures to create and verify summaries of data which can be linked to an individual mean that they provide authenticity - that the message was indeed from the original sender, and integrity - that the message was not maliciously altered. This is an essential requirement in scenarios where the identity of the sender is relevant to the content of the data, such as financial transfer requests, and is why they find several applications in cryptocurrency implementations (See section 2.3). Before discussing how each of these primitives function in cryptocurrencies in general, two other concepts require brief discussion - blockchains, and distributed computing.

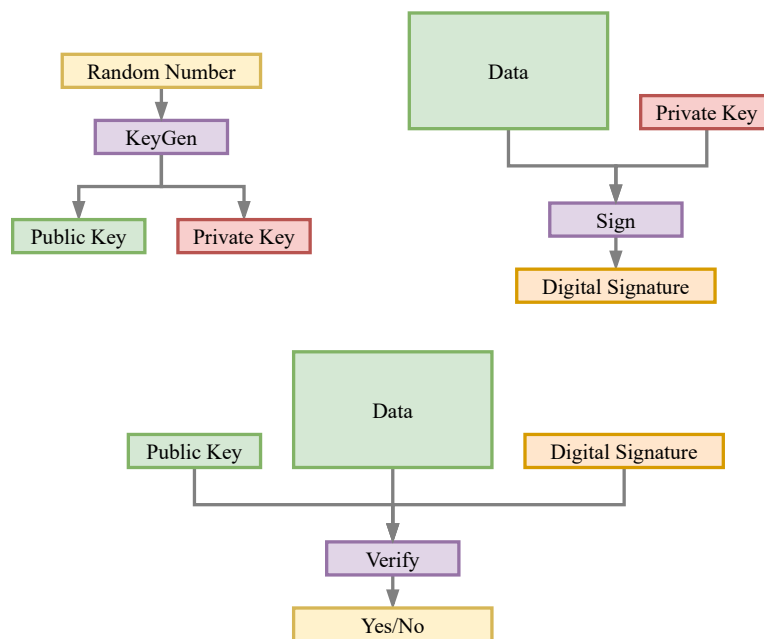


Figure 2.2: Visualisation of how the three functions of a digital signature scheme (key generation, signing function, and verification function) interact to provide authenticity to a given piece of data using a public-private key pair.




Figure 2.3: A simple blockchain with a genesis block and three appended blocks. Each appended block contains a hash of the previous block, forming a singly linked append-only chain.

2.2 Blockchains

A blockchain is a data structure in which pieces of data (known as blocks) are appended to one another to form a chain [20]. They are also stored in a distributed way across a network of computers, and use a consensus mechanism to securely append new blocks. This section describes the fundamental concepts behind blockchains, including dedicated sections to distributed computation (2.2.2), and consensus mechanisms (2.2.3).

2.2.1 Data Structure

The process of appending a block to a blockchain includes computing and storing a cryptographic hash (see Section 2.1.1) of the previous block. This cryptographic hash can be considered a unique summary, or fingerprint, of the data in that block. Since each block contains a hash of the previous block, and that block contains the hash of its previous block, and so on, this appending process creates a one-way (singly linked) chain of blocks - a blockchain (see Figure 2.3) [15]. The only exception to this process is the very first block in the chain, known as the genesis block. The genesis block can contain anything its creator sees fit, such as details about the sizes of blocks in the chain or parameters for the hashing algorithm, and can be thought of as the starting point, or root, of all of the subsequent blocks and hashes in the blockchain.

2.2.2 Distributed Computing

Computers connected to one another via the internet can communicate to complete tasks and store data, forming a distributed computing network [21]. These tasks can range from solving complex problems such as protein folding

[22], to storing pieces of data [23], to routing communications in unpredictable ways [24][25][26]. One way in which a distributed network of computers can be used is to store a blockchain. When storing a blockchain, each of the computers in the distributed network is responsible for storing a copy of the given blockchain, and can be referred to as a nodes [27]. In practical terms, each node is running an instance of an identical piece of software which contains all of the functionality needed to gather existing blocks from other nodes in the network. In addition to simply storing whatever the current state of the blockchain is, nodes need to be able to append new blocks to the blockchain as new data is added, and have these new blocks be distributed to other nodes in the network. The exact process by which all of these nodes reach a consensus on which new block to append and distribute is called a consensus mechanism [28].

2.2.3 Consensus Mechanisms

Blockchains have new blocks appended to them over time, so require some mechanism by which this can be achieved. This becomes somewhat complex as all of the computers in the network must each agree to make additions to their stored chains in a certain way, and must therefore reach some sort of consensus amongst themselves as to the content *and* order of new blocks in the chain. This way of adding new blocks to the chain is commonly referred to as a consensus mechanism, and is an important mechanism for the security of the underlying blockchain, as the introduction of new (or malicious) blocks in the chain may affect the cumulative values of previous items in the chain. The nodes on the network which participate in this consensus mechanism are called miners. An in depth discussion around different consensus mechanisms is not essential to understanding further work in this thesis, but Wang et al provide a recent (2019) survey [28] which contains a description of several consensus mechanisms used in blockchain networks for reference.

So far this chapter has described how cryptographic hashing and distributed computing can be used to create a decentralised data structure which is resistant to malicious modification. To summarise;

- A blockchain is a distributed append-only data structure
- The computers across which a blockchain is distributed are called nodes

- Nodes use consensus mechanisms to add new blocks to a blockchain

One such data type which particularly benefits from the combination of distributed storage and blockchains is financial transaction data, as this can ultimately result in a censorship resistant, globally distributed payment processing network supported by robust cryptographic principles. The following section describes exactly how this next step (from blockchains to cryptocurrencies) is made, and introduces the main topic in this thesis; cryptocurrencies.

2.3 Cryptocurrencies

Cryptocurrencies work by storing copies of a blockchain containing transaction data across a distributed network of computers [29]. This means each computer on the network is running a program which communicates with other computers in that network to synchronise and update a blockchain. As discussed in the previous section, these networks must each have a shared understanding of how the blockchain will be updated, called a consensus mechanism, and details about what data the contents of each block should be. The exact synchronisation and consensus mechanisms used vary between cryptocurrencies, and are typically unique between any two cryptocurrencies although this is not always the case [28][30].

The first cryptocurrency to function in this distributed *and* consensus driven way was Bitcoin, which was first published under the pseudonym Satoshi Nakamoto on the 31st of October 2008 [31]. However, several similar variants of what we now know as cryptocurrencies, or digital cash, had been attempted before this. Of these precursors, David Chaum's eCash - based on his cryptographic primitive of blind signatures published in 1983 [32] - is generally similar to Bitcoin in its architecture, but failed to gain widespread use following its implementation in 1995. Similarly, Wei Dai's B-money employed similar ideas in 1998 [33], again failing to gain widespread use. Finally, Nick Szabo's Bit Gold published in the same year [34] shares many similarities with all of the above, but was never implemented - although many of Szabo's ideas are identical to those presented in Satoshi Nakamoto's Bitcoin paper published a decade later [31]. These precursors to Bitcoin, and indeed the vast collection of cryptocurrencies that followed, do however highlight the pace of innovation in this new and emerging technology. A

key innovation that separates these precursors from their more successful descendants is the idea that they could store more than simply a record of transactions within their blockchains.

More advanced cryptocurrency implementations offer not only the ability to store and update transaction data, but the ability to store and execute computer code. This means that rather than using blockchains to just store data in a secure way, computer code could be executed on each of the nodes in the network, and the results of that computation (and the code) be stored within the blockchain itself. This addition enables a host of technical possibilities which combine the benefits of regular distributed computing with the benefits of robust cryptographically verifiable storage. Examples of such implementations, which can be most accurately described as distributed computing systems which also include cryptocurrency integrations, are Ethereum², EOS³, and Solana⁴.

2.3.1 Addresses & Wallets

The previous sections have described blockchain data structures, distributed computing, and how they interact to enable the creation of cryptocurrencies. However, the mechanism by which value is actually stored within cryptocurrency blockchains has not yet been discussed, but is crucial to further work in extracting the value transfers stored within these data structures. The three key components to this discussion are;

- asymmetric cryptographic keys
- cryptocurrency addresses
- cryptocurrency wallets

which represent three levels of abstraction in actually sending and receiving cryptocurrency.

Within a cryptocurrency blockchain, user accounts must have unique identifiers which can somehow be used to receive and authorise the transfer of funds. One way in which accounts can be represented is by using a public key, which can be created using a public-key, or asymmetric, cryptosystem. As

²See <https://ethereum.org/en/>, accessed 12/05/2021.

³See <https://eos.io/>, accessed 12/05/2021.

⁴See <https://solana.com/>, accessed 31/01/2022.

described in Section 2.1.2 above, asymmetric cryptosystems use pairs of cryptographic keys for different cryptographic operations, rather than symmetric cryptosystems which use identical keys [11]. In the context of cryptocurrencies this means that every user account has both a public and private key, which are simply two mathematically related numbers (typically between 32 and 256 characters long). Conceptually, these public and private keys can be thought of in the same way as email addresses and passwords, albeit with a number of useful but omitted cryptographic properties. Like emails, funds can be sent to the public key, and the transfer of funds out of the account can only be authorised by use of the private key.

As a concrete example, if Alice wanted to send some cryptocurrency⁵ to Bob, Bob would first need to generate a public-private key pair using an asymmetric cryptosystem, and then share his public key with Alice (Alice's public key may also be known to Bob thereby linking each person to their cryptographic key, this can be achieved using public key infrastructure such as a website or other out-of-band message). Alice would then create a cryptocurrency transaction and sign it (see Section 2.1.3) using her private key, which only she knows. The computers on the cryptocurrency network would then see that a new transaction to Bob had been received and that it had been signed by Alice. Since it has been signed, the nodes in the network can verify it (see Figure 2.2), and if the transaction was indeed signed by Alice then the transaction can be considered valid. They would then add this transaction to the next block in the blockchain, and whatever value Alice had specified in her transaction would now be available for Bob to spend. Once the block has been added to the blockchain (using some consensus mechanism), the transaction would appear as a transfer of an amount of cryptocurrency from Alice's public key to Bob's public key on the blockchain itself. Note that the total amount of cryptocurrency that Bob has at any given time is implied by the sum of all transactions to his public key⁶. Also note that Bob cannot maliciously create a transfer of funds from Alice, as he does not have access to her private key, and therefore cannot sign transactions from her public key (address).

In this example, the public and private keys are used for different purposes.

⁵This is a generic example, specific cryptocurrencies may use different asymmetric cryptosystems to generate keys.

⁶Although the current state of each address could be stored on the nodes in the network too.

Both Alice and Bob’s public keys can be described as their cryptocurrency addresses, which, like email addresses, can be used to send things (cryptocurrency) to them. This means that when further discussions in this thesis describe cryptocurrency addresses, it refers to the public keys which serve to uniquely identify accounts. Their private keys in this scenario act as a sort of password, allowing them to sign (see Figure 2.2) transactions which make them valid (and therefore added to the blockchain) in the eyes of each of the nodes in the network. This is much the same concept as cryptographically signing documents to verify that they are indeed from the sender. Without a private key, transactions from their respective addresses cannot be signed and therefore will not be validated by the nodes in the network, so funds cannot be transferred out of the address.

Finally, a cryptocurrency wallet is simply a device, piece of software, or simply a text file, which contains a user’s public *and* private keys. Once a user generates their public-private key pair, their wallet is the term for the location in which these values are stored. This location can be in a browser extension such as MetaMask⁷ or TrustWallet⁸, an encrypted physical device such as Trezor T⁹ or Ledger Nano¹⁰, a text file on the user’s desktop titled ‘my_keys.txt’, or even a physically written out or printed document with the cryptographic keys on. Each of these different wallet solutions offer different levels of permanence and usability. For the remainder of work in this thesis, it is assumed that cryptocurrency has been sent using a browser extension wallet as the applications in question are predominantly web applications.

To summarise;

- User accounts in cryptocurrency networks are typically represented using a public-private key pair
- The public key part of this pair is referred to as the user’s address
- The private key part of this pair is conceptually most similar to a password so is kept private by the user
- A wallet refers to the location of a user’s public-private key pair, and is required to send transactions

⁷See <https://metamask.io/>, accessed 31/01/2022.

⁸See <https://trustwallet.com/>, accessed 31/01/2022.

⁹See <https://trezor.io/>, accessed 31/01/2022.

¹⁰See <https://www.ledger.com/>, accessed 31/01/2022.

- Once verified, transactions are added to the blockchain by nodes on the network

With addresses, wallets, and a simple example transaction briefly discussed, the next important concept in understanding how to decode cryptocurrency transactions for this thesis is smart contracts.

2.3.2 Smart Contracts

Smart contracts are a type of computer mediated agreement between two or more parties over a computer network, an idea first published in 1996 by Szabo [35] and formalised one year later [36]. Smart contracts extend the idea of sending funds to each other described above by allowing either Alice, or Bob, or both, to be a computer program known as a smart contract. These contracts represent any agreement between two parties by programmatically verifying that a set of conditions was met, rather than relying on a human third party verifier as used in regular contracts. Unlike regular contracts which are written in human language, smart contracts need to be represented as a set of programmable instructions, so can be thought of as computer programs which bind two or more parties together. This use of a computer network to mediate agreements has the critical effect of removing the need for a third party, instead replacing it with the need for a computer network. Whilst implementing this idea meant that such computer networks could now exist, it was not until the conception of Bitcoin [37], and the later conception of Ethereum [27], that accessible and flexible implementations of this concept became widely known.

An important observation is that the computer network described above does not necessarily need to have payment processing capabilities built in, as a smart contract may simply return a yes/no answer to a query, for example. However, if the computer network is also a cryptocurrency network, then payments can be processed in the native cryptocurrency and delivered to the parties' cryptocurrency addresses by the contract itself. This computer mediation of financial agreements in particular - as opposed to political or social agreements for example - raises the utility of smart contracts from simply a way of verifying a set of generic conditions, to enabling the automatic transfer of wealth based on those conditions. This technology is extremely powerful and has implications across societal governance, the operation and

transparency of financial institutions, and the financial freedoms of citizens to transfer wealth in new ways, to name a few. As this thesis focuses specifically on the gambling domain, it is first important to concretely define what a smart contract is in the context of a specific cryptocurrency network and what makes this capability possible.

Ethereum Network

All of the studies in this thesis focus on data available on the Ethereum network, chosen because it is the oldest (2014) [27] and largest (approximately £240Bn)¹¹ cryptocurrency network by market capitalization which explicitly contains smart contract functionality. This means that it has a comparably mature ecosystem of applications and developer resources versus other cryptocurrencies, making it an ideal candidate for gambling research. The Ethereum network makes computer mediated financial agreements possible by combining its blockchain with a Turing complete programming language called Solidity¹² in a single abstracted layer, thereby allowing smart contracts to interact directly with its blockchain. It does this by executing a program known as the Ethereum Virtual Machine (EVM) on each of the nodes in the network, which does more than simply synchronise and update the Ethereum blockchain. Here, a smart contract is a program which governs the behaviour of accounts within the Ethereum state. These contracts are stored within the blockchain itself, so can be accessed by anyone with a computer connected to one of the nodes in the network. In plain terms, this means it can send and receive the native cryptocurrency Ether (ETH), and given Solidity's Turing completeness can express all computable tasks [27]. Such contracts can be executed by transacting with them in the same way as a regular cryptocurrency address, but with details about the function call and parameters embedded within the transaction. This mechanism is discussed in greater detail in Chapter 3. This technology therefore pulls ideas from cryptography, blockchains, distributed computing, economics, and computational theory together to create a powerful system of automatic, programmable, and secure wealth transfer which has not yet existed at such a scale in the history of human civilisation. Of course, such a powerful technology has a number of applications. One such application is gambling, and is the focus of this

¹¹At time of writing (01/02/2022), see coinmarketcap.com accessed same date.

¹²See <https://docs.soliditylang.org/>, accessed 01/02/2022.

thesis.

Smart contracts, like other pieces of computer code, can store variables which can be modified by code held in the contract itself. These variables may contain objects which relate to real world objects, events, or capabilities, and in some cases can be traded between users. In the context of gambling, these may be things like casino chips, loyalty cards, or other tokens of interest. In the context of gaming, these may be in-game items, subscription tokens, or achievements such as participation in a notable in-game event. While these variables can take a number of forms, those most relevant to this thesis are referred to as a non-fungible tokens, which is discussed in detail in the following section.

2.3.3 Non-Fungible Tokens

In economics and by extension virtual economics [38], non-fungible goods are those which are not interchangeable for equivalent value, as opposed to fungible goods which are generally similar or identical and are therefore considered to be of equivalent value. For example, the majority of food products such as apples, potatoes, and so on, are fungible in the sense that one is interchangeable for another with no change in value. Non-fungible items on the other hand are those which are unique, and which cannot readily be exchanged for another similar or even identical item. For example, the Mona Lisa - like the majority of fine art - is non-fungible, meaning that even a perfect replica would likely have a considerably different value to the original, and the two would therefore not be interchangeable.

The smart contracts described in the preceding section can be used to create tokens which can represent anything from real world goods, to virtual goods, to abstract concepts such as voting rights or access to privileged information. These tokens therefore exist simply as transactions in the blockchain, as does the underlying cryptocurrency. These transactions, and the contracts which they are defined within, follow technical standards such as those described by the Ethereum Foundation, and can be implemented in a similar way to other virtual goods; i.e. with some issuance mechanism, a transfer mechanism, and any other verification mechanisms which may be useful for proving ownership. In the context of the Ethereum network, fungible tokens can be implemented following the ERC-20¹³ standard, and

¹³See <https://ethereum.org/en/developers/docs/standards/tokens/erc-20/>, ac-




Figure 2.4: An example NFT from the Bored Ape Yacht Club collection (See <https://boredapeyachtclub.com/>, accessed 01/02/2022). This particular ape (See <https://opensea.io/assets/0xbc4ca0eda7647a8ab7c2061c2e118a18a936f13d/8585> for details, accessed 01/02/2022) sold on October 19 2021 for 696.969ETH, which at the time of writing is equal to approximately £1.4m (time of sale approximately £1.9m).

non-fungible tokens can be implemented following the ERC-721¹⁴ or similar standard¹⁵ (See Figure 2.4 for an example). In plain terms this means that for a given non-fungible token, or NFT, a smart contract exists which defines the issuance and transfer mechanisms along one of these standards, and ownership of that token can be granted by a function call to that contract.

Non-fungible tokens, and indeed fungible tokens, are two types of crypto-asset which can be created atop cryptocurrency networks, and which can be

cessed 28/04/2021.

¹⁴See <https://ethereum.org/en/developers/docs/standards/tokens/erc-721/>, accessed 28/04/2021.

¹⁵Several standards exist which vary from ERC-721 in some way but perform generally the same function.

transferred in much the same way as the cryptocurrencies with which they are defined. Because they can be transferred in this way, and because they share all of the cryptographic properties of the contracts and transactions atop the network, they draw into question their classification from an academic, legislative, and regulatory perspective [39][40]. For example, if a smart contract issues a fungible token called ‘apples’ which can be transferred freely between people and can be bought and sold for a given cryptocurrency, then this token is simply a deeper abstraction of the cryptocurrency with which it can be bought and sold for. However, the same could be said for a non-perishable fungible token in the real world such as gold. In this instance, gold ingots are not a currency per se, but rather a store of value. Understanding this distinction, and how it manifests in terms of the differences in behaviours across systems which use these types of token, is currently unknown, and beyond the scope of this thesis. However, findings presented in studies below will touch on this issue as encountered as exploration in this area is an important part of future work in cryptocurrency research.

Before moving on to discuss how NFT’s, smart contracts, cryptocurrency wallets, and the Ethereum network, can be used to create gambling and gaming platforms, the very concept of value in cryptocurrencies can be briefly discussed. This short detour will be important for framing the scale of the studies in the decentralised gambling domain presented in this thesis.

2.3.4 Price Volatility

Cryptocurrencies, as the previous sections have described, exist as distributed blockchains containing transaction data and sometimes also computer code. These transactions and code store variables which can correspond to amounts of the cryptocurrency itself, non-fungible tokens, and other variables which all exist solely within the blockchain. As these distributed computing systems exist at a global scale, their prices can be affected by changes in policy at the national and international level [41][42], price changes in other currencies [43], and fundamental factors such as the number of users and the computational power of the networks themselves [44] to name a few. The prices of cryptocurrencies are also subject to market manipulation [45][46], including coordinated pump-and-dump schemes [47] which seek to artificially inflate (and deflate) prices for a short period of time. These factors, paired with the youth of this new technology, make cryptocurrencies extremely volatile



Figure 2.5: The price of ETH (the native currency of the Ethereum network) over time, see <https://www.coingecko.com/en/coins/ethereum>, accessed 04/11/2021.

in comparison to other asset classes (see Figure 2.5). As any derivatives, such as NFTs, sit atop these systems, they too can exhibit extreme volatility, frustrating efforts to reliably value them over a given period of time.

In the context of gambling, this means that bet sizes, payouts, and the real monetary value of any tokens generated as a result of player activity can be extremely volatile. This has important ramifications when interpreting the results presented in later chapters, but also invites a number of questions around the perception of monetary value in this new and technologically advanced domain. Such questions are not the focus of this thesis, but are nevertheless related and may benefit from the findings presented herein.

Although understanding exactly which factors contribute to the market prices of given cryptocurrencies is a huge and complicated task, for gambling only the current market price is needed, as this is ultimately what will contribute to understanding behaviours in real terms. Since cryptocurrencies exist purely digitally and are not tied to a specific geographic region or country, they are traded on cryptocurrency exchanges [5] which function in the same way as stock markets but are open 24/7. These exchanges can be used to derive current and historic valuations of cryptocurrencies, which greatly simplifies mapping user actions to equivalent real monetary costs. With a brief discussion of the value of cryptocurrencies presented, one can return to the main thread of this chapter to discuss decentralised applications.

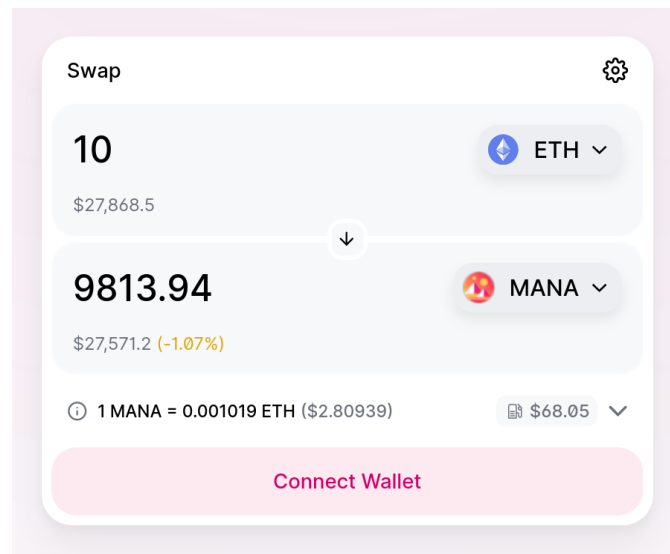


Figure 2.6: Uniswap user interface example converting from one cryptocurrency (ETH) into another (MANA).

2.4 Decentralised Applications

Decentralised applications are computer programs whose code is executed in whole or in part atop a cryptocurrency network. They typically have at least two core architectural components which operate in unison to provide both an interface for users, and the underlying computational functions needed to complete their objectives. The interfaces of these applications can be a website, mobile app, or other client-side program. The computational functions for these applications are stored partially or fully within smart contracts. These contracts are created to accept function calls from the interface, which pack the client side actions into cryptocurrency transactions and execute them across the cryptocurrency network.

An example of a decentralised application is Uniswap, a token exchange platform operating atop the Ethereum network. This application has an interface, available at <https://app.uniswap.org> (or see Figure 2.6), and a collection of smart contracts which each offer some functionality such as swapping tokens. The application's interface is simply a web page like any other, with the addition of scripts which can link the interface to a user's cryptocurrency wallet, allowing funds to be sent and received. Here, the term 'cryptocurrency wallet', rather than being simply a public-private key

pair corresponding to a given Ethereum address, refers to a web browser extension or other application which holds this information for the user (see Section 2.3.1 above). This is conceptually similar to contactless payment technologies in the physical world, which hold account information in a secure way, and allow spending without having to enter card details at every store.

Decentralised applications can take a number of forms, from offering financial services, to marketplaces, to gambling and digital games. The two most relevant to this thesis are decentralised gambling applications and decentralised gaming applications, as discussed below.

2.4.1 Gambling

Decentralised gambling applications are those which provide a gambling service of some kind such as casino games or sports betting, accepting cryptocurrency as stake [48]. The vast number of cryptocurrencies that currently exist, paired with the vast number of different ways in which they can be used, means decentralised gambling applications come in many forms. Examples of decentralised gambling applications are;

- dice2.win - a casino offering simple games such as coinflips and dice rolls
- etheroll.com - a simple 1-100 roll style application
- dicether.com - a casino offering dice, keno, and other simple games

These applications typically have simple interfaces which present a selection of games, a bet size input field, and a ‘place bet’ button or similar - similar in design to typical centralised online casinos. As this technology matures, more sophisticated decentralised gambling applications will undoubtedly become available, however in the context of this thesis only those which provide simple casino games are considered. See Chapter 7 for a complete discussion of the decentralised gambling applications studied in this thesis.

Sources of Randomness

An important detail of decentralised applications, and decentralised gambling applications in particular, is how exactly randomness is generated. Two main ways exist in the applications studied in this thesis, the first is to

use some external source of randomness and simply commit its output to the blockchain. Applications which choose this option are architecturally somewhat more complex as they must rely on a callback of some kind from the off-chain source. The second way is to use the output of some cryptographic computation which cannot be predicted or manipulated at the time a player commits to a given outcome. The mechanism by which this can be achieved relies on using the output of hash functions (see Section 2.1.1) as a source of entropy in some computation, although for additional security a commit-reveal (use hash of some value to commit to an output, then reveal the value at the point of computation) construction is also used. The reasons behind the security of this technique are somewhat technically involved and application specific so are not expanded upon here. For the purposes of this thesis, understanding the source of on-chain randomness is not essential to the analysis of transactions to a given application, but is nonetheless an important architectural influence to be aware of.

2.4.2 Gaming

Decentralised gaming applications, or blockchain games [49], are decentralised applications which provide a gaming service such as the trading card game *Gods Unchained*, or the collectable kitten breeding game *Cryptokitties* [40]. These games differ from the decentralised gambling applications described above in that they typically enable the generation and use of some virtual good, such as trading cards or virtual kittens in the case of these two particular blockchain games. In the same way that the landscape of centralised digital games differs dramatically from centralised casino games, so too does the range of mechanisms within blockchain games vary dramatically from those found in decentralised gambling applications. However, as in other forms of digital game, blockchain games can incorporate chance based mechanisms into their design, making them mechanically similar to simple casino games [9], and therefore of interest in this thesis.

Like their centralised counterparts, decentralised gaming applications are extremely diverse in their architectures, genres, and mechanics [49]. As such, a complete discussion of all of them is beyond the scope of this thesis. Instead, see Chapter 6 for an in depth discussion of those studied at the aggregate level, and Chapter 8 for a specific case study and analysis of the behavioural distributions of players which use a mechanism within it.

2.5 Summary

This chapter has presented the fundamental building blocks needed to understand how cryptocurrency networks work, and how they can be used to create decentralised applications. It has shown how blockchain data structures can be distributed across a computer network to store financial transactions, and how these networks can make use of modern cryptographic primitives to form robust, secure, payment systems capable of transacting value at a global scale. The issue remains of how this understanding of the fundamental building blocks translates into actually gathering and decoding cryptocurrency transactions for academic research. The next chapter is therefore dedicated to understanding how this theory can be applied in the real world to extract actual cryptocurrency transactions, decode their smart contract function calls, and identify application-specific mechanisms in decentralised digital and casino games.

Chapter 3

Decoding Cryptocurrency Transactions

“Meesa not a Jedi.”

Darth Jar Jar Binks
The Clone Wars

This chapter includes concepts and rewrites of work published in CHI Play titled ‘*Ethereum crypto-games: Mechanics, prevalence, and gambling similarities*’, and work published in PLOS One titled ‘*Inside the decentralised casino: A longitudinal study of actual cryptocurrency gambling transactions*’.

As discussed in the previous chapter, behavioural data such as bet placements, NFT minting, and other financial transactions are all stored within cryptocurrency blockchains. In order to understand the use of cryptocurrency technology, this data must be retrieved, and variables from these transactions extracted to understand which user actions they represent. This chapter begins with Section 3.1, which describes in practical terms how data stored in cryptocurrency blockchains can be accessed. Section 3.2 then describes the two different types of transaction stored within the Ethereum blockchain. Section 3.3 then describes the process for identifying function calls within these transactions, an essential process for transforming the encoded data on the blockchain into a usable academic data set. Section 3.4 then details how in complex applications, part or all of the internal mechanisms of the application require emulation in order to compute run-time variables from

transactions, rather than simply extracting the transaction parameters alone. Finally, Section 3.5 describes the process of matching pairs of transactions, where the second transaction represents the outcome of whichever function was called in the first, e.g. the payout of a bet.

3.1 Accessing Cryptocurrency Transactions

Chapter 2 described the fundamental building blocks of cryptocurrencies, and showed that they store transaction data in synchronised blockchains distributed across a network of computers. Two methods exist for accessing the transaction data stored within such cryptocurrency blockchains. These are node synchronisation, and the use of public API's. This section outlines and discusses each through the lens of research accessibility.

3.1.1 Node Synchronisation

Node synchronisation refers to the process of executing a cryptocurrency's synchronisation program on a local machine. For cryptocurrencies with public blockchains this program is typically available through a public repository provider such as GitHub. For example, one can freely access the program executed by nodes on the Ethereum network via the Ethereum foundation's github page¹. When executed, such programs connect to other devices in the network and begin downloading and verifying the entire blockchain of that cryptocurrency to local storage from its peers. Nodes can be configured in different ways depending on their available hardware and usage requirements². In the Ethereum network, three types of node can be configured;

- Full node - stores full blockchain data, participates in block validation, provides data on request
- Light node - stores header chain data, can verify data against these headers, typically used in low-capacity devices
- Archive node - a full node which builds an archive of historical states, allowing full historical access to all events in the blockchain

¹See <https://github.com/ethereum/go-ethereum>, accessed 03/02/2022.

²See <https://ethereum.org/en/developers/docs/nodes-and-clients/#node-types> for a complete description, accessed 03/02/2022.

For academic research an archive node would be most useful, however in the case of the Ethereum network an archive node requires several terabytes of storage³. Furthermore, in both full and archive nodes each of the blocks in the downloaded blockchain need to be executed locally, i.e. every action ever performed on the blockchain must be verified by your local machine. Once verified, transactions can be retrieved by querying the synchronisation program as if it were an API, the exact data returned will depend on the type of query executed. It should be noted that light node synchronisation is not useful for accessing historical transaction data as they do not store the historical state of the Ethereum virtual machine, and so cannot be used to retrieve all transactions to/from a given application, for example.

This full blockchain verification would not be a problem were the underlying data structure highly parallelisable and therefore efficient execution mechanisms were available, however, as the underlying data structure is a blockchain, by definition the blocks must be verified in series. This introduces a significant time overhead of several weeks or months for full or archival node synchronisation to occur even on advanced CPU architectures using solid state storage. For these reasons - cost and time - the node synchronisation method for accessing cryptocurrency transactions should only be chosen should the physical resources be available to perform it (high capacity SSD storage and a powerful CPU, with a stable and high capacity internet connection), else public API's are preferred.

3.1.2 Public API's

The high financial and time costs of node synchronisation make it prohibitively expensive in the context of typical research, but fortunately several public API's exist whose fully synchronised and archival nodes are available via web requests. In some cases these are simply nodes of the cryptocurrency network itself configured to accept data requests from anywhere, subject to their own terms of use, but public API's may also take the form of dedicated data providers for developers and/or academics for the explicit purpose of analysing blockchain data. In the case of the Ethereum network, the independent block explorer and analytics platform Etherscan⁴ provides

³See <https://etherscan.io/chartsync/chainarchive> for up to date chain size, accessed 23/04/2021

⁴See <https://etherscan.io>, accessed 24/04/2021.




Figure 3.1: A simplified representation of the difference between normal and internal transactions, with a human user and two smart contracts. Both the human user and smart contracts have cryptocurrency addresses.

Ethereum developers with an API as a community service. This API can be used to access all historical transaction data from the Ethereum blockchain, subject to a rate limit, without having to fully synchronise a node locally. It therefore dramatically reduces the data gathering time for small to moderately sized experiments, and is free to use. The disadvantage of such API's is that the data gathering process is dependent on the availability of the API.

All of the data gathered in connection with this thesis has been gathered through the Etherscan API, although access to a second API provider ArchiveNode⁵ was sought and granted as insurance against any changes or outage of the former. ArchiveNode operates a request only access model, so for future work should the Etherscan API become unavailable one should request access.

3.2 Transaction Types

The transactions gathered via either of the methods discussed in Section 3.1 above will typically arrive as JSON objects representing a collection of transactions for a given cryptocurrency address or block (depending on API call/request made⁶). The simplest way to save this data is in folder for each address, although in practice extracting only the data specific to the study (such as sender, recipient, amount, etc) in question should be pruned and then saved in order to avoid bloated memory usage (such as storing the hash values for each transaction, for example). It is important to note that there are two types of transaction which this section summarises, although as I discuss below this naming convention is less than intuitive.

⁵See <https://archivenode.io/>, accessed 24/04/2021.

⁶See <https://docs.etherscan.io/api-endpoints/accounts#get-a-list-of-normal-transactions-by-address> for an example API call, accessed 07/02/2022.

3.2.1 Normal Transactions

A ‘normal’ function call, as opposed to an ‘internal’ function call, can be described as any transaction originating from an address whose origin is not a smart contract (see Figure 3.1). The official Ethereum documentation describes this as ‘an action initiated by an externally-owned account’⁷ i.e. an account managed by a human rather than a smart contract. For example, a user purchasing some Ether through an exchange, then using that Ether to place a bet, would typically generate a normal transaction to a smart contract which would call some gambling function. This would appear as a normal transaction in the gathered data and could be analysed accordingly. The process of identifying these calls is described in more detail in Section 3.3 below.

The majority of transactions of interest in association with this thesis will be ‘normal’ transactions as they intuitively represent the direct actions of a human. However, not all transactions in the Ethereum blockchain are ‘normal’, and are of broader interest in terms of understanding the resulting effects of a human’s actions. For example, one ‘normal’ transaction may represent the action of a user, but that action may trigger a smart contract to perform an action which itself would be another (internal) transaction on the blockchain. This issue of user action subroutines is expanded upon in an analysis of the blockchain game CryptoKitties in Section 8.2.2 later in this thesis.

3.2.2 Internal Transactions

In the case that a smart contract itself transfers value to another address, such as paying out a bet or generating a virtual good, the transaction will be ‘internal’. Here, the term ‘internal transaction’ is not completely accurate, as value transfers out of a smart contract are not explicitly transactions per se, but rather changes in the state of the Ethereum virtual machine. Nevertheless, it helps to conceptualise ‘normal’ and ‘internal’ transactions as incoming and outgoing value transfers from the perspective of a given smart contract. An example of an internal transaction would be a decentralised gambling application smart contract paying out an amount of cryptocurrency

⁷See <https://ethereum.org/en/developers/docs/transactions/>, accessed 24/04/2021

Source Code	Encoded Contract
<pre>function make_diagram(size, color){ diagram d; d = new Diagram(size); d.color = color; return d; }</pre>	<pre>66756E6374696F6E206D616B655F6 469616772616D2873697A652C2063 6F6C6F72297B0D0A2020202064696 16772616D20643B0D0A2020202064 203D206E6577204469616772616D2 873697A65293B0D0A20202020642E 636F6C6F72203D20636F6C6F7...</pre>

Figure 3.2: Example of a basic function whose source code (in green) is compiled and encoded (in yellow) when stored on the blockchain. It is clear that in cases where the source code for a given contract is unavailable, determining its functions and capabilities becomes a reverse engineering problem.

to a user following a winning bet.

3.3 Identifying Function Calls

The existence of smart contracts in some cryptocurrency networks means that transactions to these contracts typically represent function calls. In the context of this thesis, these function calls may be `place_bet()`, `settle_bet()`, etc. Identifying which functions have been called in a given contract is largely cryptocurrency specific, as different cryptocurrencies may use different conventions, hashing algorithms, etc. As this thesis focuses on the Ethereum network, this section is dedicated to concepts specific to smart contracts on the Ethereum network, although some concepts may apply to similar networks.

3.3.1 Solidity Programming

All smart contracts atop the Ethereum network are written in the Solidity programming language [27]. This language is similar in structure to Python and JavaScript, and targets the Ethereum Virtual Machine which runs on all nodes in the Ethereum network. When a smart contract is deployed, its source code is compiled down and stored within the blockchain in an encoded form (see Figure 3.2). This encoded form is not self-describing so is specified in the Solidity documentation and expanded upon below.

The Solidity language, as with other modern computing languages, contains many different structures and capabilities which are beyond the required understanding for this thesis - e.g. multiple inheritance and interfaces. The key concept for decoding transactions in the gambling and gaming domains are the use of method signatures (detailed in Section 3.3.3), which are found in each function call to a smart contract and can be decoded to provide the exact function being executed by that transaction, along with the parameter values therein. The terms ‘method’ and ‘function’ here are used interchangeably.

3.3.2 Gas Fees

In the Ethereum network, gas is a unit of measurement for the amount of computational power needed to execute a given operation or set of operations. Before moving to discuss the role of method signatures and parameter counts in identifying function calls, a brief overview of gas fees is required to set the scene for understanding why in some cases transactions may be harder to analyse than expected. As defined by Ethereum’s yellow paper [50] (Appendix G: Fee Schedule), the minimum amount of gas required for any transaction to succeed is 21,000, plus any other costs which result from operations relating to that transaction. This means that for a transaction to succeed, users must pay a certain amount of gas to the miners (see Section 2.2.3) whose computers constitute the Ethereum network. For example, simply sending Ether from one address to another could be completed with exactly 21,000 gas attached to the transaction. However, asking a smart contract to perform some complicated computation may take hundreds of thousands of gas depending on its complexity [51].

A second consideration when associating gas with a transaction is the price of that gas, which is a parameter set in the transaction used by miners to determine whether or not to accept and process a transaction. Conceptually, this can be solidified in the following scenario; I’d like a courier to complete a delivery, and want to give them an incentive do do so. I give them a bag of 100 coins which they may use to complete tasks on their journey - a train ticket may be 20 coins, a ferry trip may be 50 coins, etc. The number of coins used per task is not set by either of us, and is a property of the world we live in (i.e. the network). I also say that I’m happy to pay 10 pence per coin used, so if all of the tasks on their journey cost a total of 80 coins

they will earn £8. My arch rival wants to complete a similar delivery, but offers 90 coins and is happy to pay 8 pence per coin. The courier can still complete their delivery as all of the tasks still take 80 coins, leaving 10 spare, but the courier will only receive £6.40 if they take my rival's offer (because the number of coins actually used is the same, but I am willing to pay more per coin). This is a simplified example of how miners on the Ethereum network operate, accepting and rejecting offers, processing or not processing transactions and executing functions, based on the amount they stand to gain by doing so.

$$cost = gasQuantity \times gasPrice \quad (3.1)$$

The amount of gas attached to a given transaction, and its price, is set at the point of payment, past which if the amount of gas provided is too low for the computation to complete, or the smart contract rejects the transaction for some other reason, the transaction fails. A failed transaction will still appear in the blockchain with a 'failed' status, the gas originally used will be lost, and any transfer of funds will be reverted. Alternatively, a successful transaction will use no more than the total amount of gas associated with that transaction, and will refund the rest to the original account (the courier will return any unused coins). This can happen when a user wants a transaction to execute quickly so would pay a higher gas price (not gas quantity, see Equation 3.1) - a miner on the network is then more likely to accept the transaction for processing and will return any gas (quantity) that is not used. In the context of gambling apps, gas fees and gas prices are not necessarily relevant to the analysis as they will be broadly similar for all players at a given time. However, in the case that gas prices be dramatically different between time periods (i.e. the cost to play a particular game or interact with a particular smart contract is greater or lesser), it may be meaningful to factor in gas prices into the analyses. This holds across different cryptocurrency networks which use a similar payment incentive for processing on the network.

3.3.3 Method Signatures

Like in other programming languages, Solidity's method signatures represent the name of the function being called, and the values of any parameters




Figure 3.3: The structure of an Ethereum transaction’s input field on a method with one parameter.




Figure 3.4: Example data set of transaction input fields showing that the function being called (e.g. `place_bet()`) can be identified by the signature at the start of the input field.

it accepts e.g. `make_coffee(container, amount, milk, sugar)` [50]. These parameters can be of different data types, e.g. `enum` for container, unsigned integer for amount, etc, and each of these parameter data types can be used to create unique signatures - even with overloaded functions. In the case of the Ethereum network, transactions can not only transfer cryptocurrency, but can also execute methods in smart contracts. This is done by encoding the parameters and appending them to the method signature of the method the transaction is calling. In plain terms, an Ethereum transaction can have not only a sender, receiver, and value, but also an ‘input’ containing a method signature and some parameter values.

Understanding this mechanism by which transactions can be used to call methods is central to decoding their contents, and ultimately understanding which transactions are doing what on the network. To apply this in the context of decoding cryptocurrency transactions, the ‘input’ field of each transaction can be decomposed into the method signature and the parameter set (see Figure 3.3). Following the Solidity documentation on transaction encoding, the method signature is always ‘the first four bytes of the Keccak hash (see Section 2.1.1) of the ASCII form of the signa-

ture'⁸. Following the example above, the ASCII representation is simply 'make_coffee(uint8,uint32,bool,bool)'⁹, the Keccak-256 hash of which is 'bffee939cb6f01ce9665e4...'. Taking the first four bytes gives us 'bffee939', so in transactions to a smart contract with a make_coffee method as above, any of those with input values which begin with 'bffee939' are calls of that function (see Figure 3.4).

It is important to note that this process of method signature matching is only possible when the source code for the contract is available, as it relies on being able to compute the signature yourself and then looking for matches in the data. This is of course possible by simply looking at the unique method signatures of transactions to a given smart contract, but lacking the context of the source code it may be difficult to determine exactly what is happening. An alternative but generally less accurate means of identifying function calls is by simply counting the number of parameters, under the assumption that any contract of interest has few methods in the first place thereby reducing the likelihood of collisions. Understanding how parameters are encoded is also useful for extracting the values passed to the method, which may be things such as the game being played, the players choice of outcome, and so on.

3.3.4 Parameter Counts & Values

Following the method signature in any transaction input field is an encoded list of parameters, which are the values passed to that method. Each parameter exists as a padded 32 byte chunk or multiple 32 byte chunks as shown in Figure 3.3, which can be extracted using a simple split operation on the remaining part of the input once the method signature has been removed. The parameter count then is simply the length of the list of chunks. For each parameter in a given method, the actual value of the parameter is simply the decoded value of the chunk extracted directly from the transaction input field. The data type of a given method can be retrieved if the source code of the smart contract is known. In practice, part or all of the source code for a given gambling or gaming DApp is typically published by the developer in

⁸See <https://docs.soliditylang.org/en/develop/abi-spec.html#examples>, accessed 22/04/2021.

⁹The enum type is mapped to an 8 bit unsigned int in the version of Solidity (0.8.5) specific to work in this thesis.

the interest of transparency.

3.4 Emulating Application Mechanics

In a simple application, the parameters and name of a function being called are all that are needed to begin detailed analysis of user behaviours. However, the gas fee system inherent to the Ethereum network acts as an invisible hand, pushing application developers toward computationally efficient and highly compressed implementations of simple tasks. This is both to lower the costs to their users of interacting with their applications, and to lower their own expenses as any transactions that their smart contracts attempt (e.g. to their other contracts such as a treasury or croupier) must also pay a gas fee. Following the example used in Section 3.3.2 above, this is akin to the courier working to make the delivery route as simple and cheap as possible so that they can offer their service for fewer coins and thereby undercut their competition.

This aggressive simplification does not always frustrate analyses as it can force developers to remove any excess or inefficient code, however, in some cases it can lead to solutions which require emulation (at least once, if not per transaction) to extract useful data. An example of such an implementation is the dice2.win platform's placeBet function, which will be the focus of a full subsection in Chapter 7. In this case, the function performs a generic chance based calculation based on modulo and betmask values which determine the type of game being played, and the user's outcome selection respectively. This is very useful from the application's perspective as multiple games can be implemented via calls to the generic chance function with different modulo values, but it also means that in order to extract the desired outcome chosen by the player, the method's internal mechanics must be emulated. In other words, one can tell from the function call what the game being played is because the modulo value will be fixed, e.g. 6 for a single dice roll, 36 for two dice, etc. Yet the betmask value represents the player's choice before an application-specific calculation is made, therefore in order to extract the derived 'rollUnder' variable (the value a random roll would need to be under in order to win), the internals of the method need to be at least partly emulated.

This process of extracting application-specific variables in order to better

understand how the choices players are making maps to low level transaction data can be very time intensive for complex applications. This complexity is compounded in the case of crypto-games, which with their multiple currency layers, items, and other features may draw on many methods in many contracts. This hard mental resource limitation is discussed as encountered in Chapters 7 and 8, and has a strong effect on the types of behavioural data available. For this reason it contributes to the choice of application for study, and the type of analysis performed. With the potential for emulating application mechanics discussed, the next important step in decoding cryptocurrency transactions is matching the bets to the payouts. This is especially important in gambling applications but can be abstracted into the digital games domain via item rarity mapping, a concept discussed in detail in Chapter 8.

3.5 Matching Payouts

Behavioural measures in the loss domain, described in Section 5.2.3, require not only the transactions corresponding to the placement of a bet, but also the payouts received by the player as a result of those bets. Whilst at the highest level of aggregation (entire career) only the sums of bets and payouts are required to compute these measures, more granular measures require matching individual bets to individual payouts. This matching process first appears trivial, however, the use of cryptocurrency networks for processing payments introduces a variable length time delay between bet placement and payout transactions. This delay means that multiple bets placed in rapid succession may result in payouts being processed out of order. Two possible solutions I propose to this payout-ordering problem are the naive chronological stitch method, and the breadcrumb trace method, outlined in the subsections below. These solutions are posed as my understanding of the data returned by the Etherscan API is that it does not return the stack trace created when an internal transaction is made (which would allow perfect matching). This means that a deeper understanding of this API or indeed the mechanisms within the Ethereum network may render this section obsolete, but in the context of understanding gambling transactions it is still useful.

3.5.1 Naive Chronological Stitch

As the name implies, performing a naive chronological stitch between a set of bet placement transaction and payout transactions means ordering them in time and mapping them one-to-one in their chronological order. This follows the overly optimistic assumption that each transaction has been processed in uniform time, and that no two transactions ever overlap.

One advantage of employing this matching system is that it is trivial to implement, however, the underlying assumption is catastrophically broken not only when a large and small bets overlap - such that a small bet appears to have an impossibly large payout - but when any of the bet or payout transactions fail. One transaction failure when applying the naive chronological stitch method can offset an entire side by one step, rendering any more granular derivative analysis useless. Additionally, employment of such a brute force approach implies zero knowledge of the structure of the application in question. For example, it may be that all payout methods contain a cryptographic hash of the bet method, or some similar nonce which ties each transaction pair together. Operationalising this application specific knowledge means applying the more sophisticated breadcrumb trace method described below.

3.5.2 Breadcrumb Trace

An alternative to the naive chronological stitch outlined above is what can best be described as a breadcrumb trace method. This method involves a manual exploration of the parameters and architectures of bet and payout methods within a smart contract, with the goal of identifying a shared value between them. A shared value, or nonce, generated by the first function and used by the second allows exact pairwise matching for each bet-payout pair. This method, unlike the naive chronological stitch above, is less sensitive to transaction failures and other ordering problems, and is more accurate for all but the most temporally distant bet-payout sequences. Furthermore, it can generate perfect data in the sense that every bet and payout is accounted for, allowing granular explorations of player behaviours using the most detailed behavioural measures.

A significant disadvantage of this technique however, is that it can be extremely human-resource intensive in the form of studying each of the

functions in a given smart contract. This is because in order to apply this technique at least some understanding of the execution process of a given application is required, although once gained this could be automated. This means that when applying a broad analytical stroke to understanding spending in the decentralised application domain e.g. multiple application's transactions, the naive chronological stitch - despite its disadvantages - may be the only option. These challenges are addressed in detail when they are encountered in practice in the following chapters.

3.6 Summary

This chapter has described the core technical principles required to begin decoding cryptocurrency transactions. As the introduction outlined, blockchain data structures and distributed computing can be combined with the addition of cryptographic primitives to form cryptocurrency networks. These networks can not only be used to transfer value, but in some cases can execute smart contracts which form the foundation of decentralised digital and casino games.

Accessing these transactions can be somewhat resource intensive if a public API is not available, however for the use cases in the following chapters API's are available. Furthermore, normal and internal transaction types make distinguishing between bets and payouts, or purchases and rewards, simpler, but still more is needed before meaningful data can be extracted. Function calls can be identified using transaction's input parameters, and their method name and parameters extracted thereafter. Source code availability makes this process much simpler, but cannot be relied upon for all decentralised applications. In some cases, identifying the type of transaction, the function it calls, and the parameter values it uses, is still not enough, in which case emulating part of the application itself may be required. This is an extremely intensive process and is practically impossible to do for all decentralised applications combined as there are simply too many, and they can be extremely complex. Finally, with all of the above information extracted, there is still one more step - matching bets to payouts. In the case of casino games this is an often simpler task than for digital games given that bets and payouts are typically paid in the same currency, but as discussed in Chapter 8 techniques can be applied to transfer these ideas into

the randomised reward mechanism domain.

Chapter 4

Player Behaviour Tracking

“An algorithm must be seen to be believed.”

Donald Knuth

So far, this thesis has described the technical concepts underpinning cryptocurrencies, including their basic cryptographic building blocks and how to extract data from their blockchains. This chapter introduces the second field of background literature needed to address the research questions posed in Chapter 1: player behaviour tracking.

A small portion of players who engage in gambling experience some form of gambling related harm [8]. An area of gambling research known as player behaviour tracking aims to understand how players within a population engage in different types of gambling, and how different engagement patterns across these types of gambling relate to experiencing harms [1]. One way in which engagement in gambling can be modelled is through the use of transaction data. That is, data corresponding to the placement of individual bets and payouts. Behavioural profiles can be constructed using this data, and these profiles correlated with external harm related variables such as self reported problem gambling severity [52] or self exclusion due to problem gambling [53].

This chapter reviews existing literature on the use of these behavioural measures to understand player behaviour. It begins with Section 4.1 providing an historical overview of the types of gambling, followed by 4.2 describing types of games in which transaction data are generated, then an overview of the structure of the data they provide in Section 4.3. With basic gambling

games and other activities described, Section 4.4 then presents the results of a systematic literature review, the aim of which was to gather the collection of behavioural measures used in player tracking research, so that the current collection of state of the art measures can be extracted. Given the importance of understanding the behavioural measures used across existing literature, a detailed discussion of each of their computations and motivations is deferred to a dedicated Chapter 5. Instead, Section 4.4 describes how broadly the literature applies player behaviour tracking techniques, and a high level summary of the breadth and depth of behavioural measures used.

4.1 Types of Gambling

Gambling has formed a part of human activity since the dawn of civilisation, or as Schwartz writes; is ‘simply older than history’ [54]. This long history means a vast number of different forms of gambling currently exist, ranging from supernaturally attributed acts of divination, through to blood sports, team sports betting, mechanical games, card games, political betting, and more. These many branches each integrate two key concepts, namely some chance based event, and payment of some kind based on that event. More concretely, Griffiths poses four criteria to be met for an activity to be considered gambling in his discussion around loot boxes in digital games [55], derived from his earlier set of five criteria [56]. These criteria are;

- The exchange is determined by the outcome of a future event, which is unknown at the time of betting
- The outcome of the future event is at least partly due to chance
- An exchange of money/objects of financial value occurs, typically without productive work from either side
- Losses can be avoided by not taking part in the activity
- (earlier & removed) Winners gain at the sole expense of losers

Each of the gambling applications, and indeed parts of the gaming applications, studied in this thesis meet all four of these criteria, as do the activities studied in the field of player tracking described above. It is however important to note that this set of criteria, whilst widely used in the literature,

are a purely academic definition. In the case of the relationships between gambling and gaming it is also important when framing findings to consider relevant legal definitions around these topics. That said, detailed discussion of legal definitions in different jurisdictions and across different activities is out of scope of this thesis, but will be briefly discussed as encountered in discussions of results below.

4.1.1 Mercantile versus Social Gambling

Given many different types of gambling activities exist, which can be configured for play at different speeds, a third important descriptor of such activities is the type of party from which the payouts are paid. Mercantile gambling, also known as commercial gambling, encompasses all games in which a player competes against an institution or ‘house’ [54]. This means that it is the gambling operator (merchant)’s role to provide the mechanisms by which the player can gamble, ensure that these mechanisms are fair, and that payouts are delivered. Different jurisdictions around the world have different regulations which describe exactly how operators can provide these mechanisms, the nature of the specific mechanisms, and which licenses are required in order for the operator to provide access to the mechanisms to their customers. Typical examples of mercantile gambling games include Roulette, Baccarat, and all manner of electronic or physical terminal type games such as video poker and slot machines which can typically be found (although not exclusively) in casinos.

Social gambling encompasses all games which are played between two or more people [54]. Although social gambling games themselves may be mediated by a third party, the payouts are directly funded by other players. This detail is captured in Griffiths removed criteria for gambling activities that winners gain at the sole expense of losers - the removal of this criteria means that his revised set can include both mercantile and social gambling. This lack of a ‘house’ in the mercantile sense means that the third party mediating the game may take a portion of the bets per round as payment for their mediation services. Any ‘house’ edge in social games therefore comparable to the house edge in mercantile gambling in the sense that it represents money lost by all players - a feature which will become important when later discussing the computation of behavioural measures. Examples of social gambling includes Poker, (social) Baccarat, and Mahjong. As with mercantile

gambling, social gambling can be created around essentially any chance based mechanic, although following the taxonomy outlined by Gainsbury et al [57], should the outcome be determined by skill alone the activity can be classified as a tournament or competition rather than gambling.

It should be noted that there are branches of gambling and games studies which focus on so called ‘social casino games’ [58, 57]. Social casino games are typically implementations of casino games (both social *and* mercantile), which instead of using a government backed (fiat) currency as stake instead use a token of some kind which is specific to that game or user account. These tokens typically do not hold any real world value, and cannot be exchanged for real world currency at the game level¹, but nonetheless function as scrip in the context of these games. Social casino games are not simple casino games - the focus of this thesis - so are not discussed further.

4.2 Types of Game

Of the many different forms of gambling discussed above, the most relevant to this thesis are gambling games, specifically casino games - as opposed to sports betting etc. Casino games come in many forms, and central to distinguishing between different types of game within this sub-category is the role that skill and chance play in the outcomes. This section therefore discusses some of the many different types of game available, and specifies which on this broad spectrum this thesis is most concerned with.

4.2.1 Games of Skill

Games of skill include all games in which the outcome is not wholly probabilistic, in other words, more skillful players will usually beat less skilled players. Examples of games of skill include Chess, Backgammon, and Mahjong. Such games can become gambling games when wagers are placed on the outcomes, or are integrated into the rounds or specific player actions within the games themselves. For example, a game of chess can be crudely converted into a gambling game by placing a payout penalty of £10 on each of the players rooks. In this scenario, a player that takes two rooks and wins the game would receive a payout of £20 should they have both their rooks remaining.

¹As with all virtual economies, trading at the account level i.e. selling and buying account with virtual currency inside, can still occur.

In other words, an exchange of money can occur based on some in-game outcome. As one can imagine, simply adding financial penalties to games of skill can drastically change the very objectives of the games, so is uncommon in gambling games in general, as purpose-built games have been created which better capture financial commitments.

One such game of skill in which financial commitments are integrated into the game itself is poker. Interestingly, the debate around whether or not poker is a game of skill or chance is still openly contested, with legal cases in the USA ruling that poker (Texas Hold'em) is predominantly determined by skill [59]. However, this does not consider the duration of play, as a single hand in poker between two players is almost completely probabilistic, yet playing many hands allows the elements of skill in the game e.g. bluffing, stake management, and folding strategies, to play a more substantive role in the long term outcome. Indeed, the very existence of professional poker players indicates the innate ability of some players to become successful enough over a sustained duration to form a career from their play - although whether this career can be attributed to tournament winnings or simply sponsorships is another question. Skill and chance are not mutually exclusive, and coexist in many gambling games. Games of skill are not the focus of this thesis so are not discussed further.

4.2.2 Games of Chance

Games of chance include all games in which the outcome is purely probabilistic, where player skill has no effect [60]. Games of chance include Roulette, Craps, and lottery type games including Bingo. These games are extremely prevalent, and are essentially ubiquitous in physical and online casinos globally. Importantly, from an analytical perspective, individual skill in games of chance cannot be evaluated as it simply doesn't exist. For example, analysis of a Poker player's betting patterns may reveal details about their knowledge of the statistical probabilities of winning their hands, indicating a level of skill in their play. However, no such analysis is possible on the betting pattern of a Roulette player as no player actions can meaningfully influence the outcome. One caveat to this difference is that betting systems can be used across both chance based and skill based games, which effectively leverage a player's bank roll against the probabilities inherent to the game [54]. An example of such a system is the Martingale (double-or-nothing) system, whereby a

player doubles the size of their bet if the previous outcome was a loss. This ensures that given a large enough bank roll, and a short enough sequence of losses, a player will always recover their losses. Of course, this strategy is inherently flawed as just six losses in a row with an initial bet of £10 would require a single £1,000,000 bet, but nevertheless the existence of betting systems as tools for players of chance based games should be acknowledged.

4.2.3 The Skill-Chance Spectrum

Many of the games discussed above, despite being grouped into their skill/chance subsections, can be more accurately described as existing on a spectrum with skill on one end and chance on the other. For example, chess would exist at the purely skill end, and a coin flip would exist on the purely chance end. As discussed in the case of poker, duration of play - in addition to just the mechanics - can play a role in the amount of chance or skill that a given gambling activity requires. All of the studies in connection with this thesis focus on games which exist solely at the chance based end of this spectrum. This is because they are typically simpler to analyse, as methods from coin flip games can be directly applied to dice rolls, for example, and because they are the type of game which currently available implementations of cryptocurrency technology allow. A brief discussion of the types of data generated by these games is now required, which is then followed by a systematic review of their academic analyses.

4.3 Data Types

The types of game described in section 4.2 each generate behavioural data, yet this data can be captured in different forms and at different levels. These include account level data, and transaction level data, which can be used to explore behaviours at different levels in the player-game interaction. This section briefly introduces the different types of data so that in later sections these variables can be the focus of specific investigations.

4.3.1 Account Level

Account level transaction data concerns incoming and outgoing funds with respect to a player's account, e.g. deposits and withdrawals. This naturally

only applies to operators and activities where an account is required. This high level data is useful for comparing non-granular behaviours between players, but is limited in that it cannot be used to understand behaviours within a game itself, only the player's external financial interactions with the operator. Many studies employ account level data to understand gambling behaviours in the context of income data, geographic data, and other meta-level variables which assess the social and economic impacts of gambling at the highest level. As this thesis focuses on the use of transaction data specifically, further discussion around account level data and derivative analyses is not explicitly relevant. That considered, many account level studies appear in later sections and provide useful insights for framing findings based on more granular data.

4.3.2 Transaction Level

Individual transaction level data is the most granular data possible regarding in-game payments. In gambling terms, each observation corresponds to the placement of a single bet, therefore capturing all in-game activity in casino games, and any financial activity in digital games. Transaction level data can therefore be used to generate detailed player behaviour profiles depending on the measures used, but are also subject to significant limitations which are discussed in subsequent chapters as encountered. Importantly, all gambling transactions have a fixed set of variables, or anatomy, upon which all behavioural analyses are built.

Anatomy of a Bet

The simple casino games most relevant to this thesis have a very basic anatomy that generates a fixed number of variables for each bet placed. Mapping these variables is essential to understanding which behavioural measures can be derived from their sets, and how these measures actually represent the behaviours and actions of the player. Taking the example of a simple coin flip game, four key variables exist which can completely describe the placement of a bet. These are (i) the player's identity, (ii) the size of the bet, (iii) the time the bet is placed, and (iv) the odds of the bet. The data types of the first three values can be simply integers (taking a unix time representation), however the odds of a bet can be represented in several

ways. For the analyses performed in association with this thesis, all odds are shown using decimal point representation, which in the case of a coin flip would be 2.0. This means that for every £1.00 bet, upon winning a player would receive £2.00. Alternative representation schemes include fractional representation and US or moneyline representations. Decimal odds are most intuitive to interpret so are the only ones discussed here.

From these four values, derivatives such as the expected value of the bet (£1.00), and the maximum payout can be computed. Additionally, in the context of simple casino games, the type of game being played can be derived from the odds as there are a finite number of games each with a fixed number of possible outcomes. For example, odds of 2.0 represent a coin flip, odds of 6.0 represent a single dice roll, odds of 36.0 represent a two-dice roll (for doubles), and so on. This derivation of game type from odds breaks down for more complex games, but holds for the simple casino games explored in this thesis. Any work exploring gambling and gaming transactions through the lens of gambling studies must therefore meet these minimum data requirements for each bet, or provide aggregate values thereof. As the following systematic review reveals, granular transaction level data has been historically difficult to acquire [1], yet in the context of cryptocurrencies this level of data access is the norm, albeit in an encoded form (as demonstrated in Chapter 3).

Before moving to a systematic review of the use of transaction data in existing research, a note on payout data is needed. Specifically, payout data is the natural mirror of bet data, so contains all of the same variables minus the odds. From a behavioural analysis perspective the bet placement transaction data is therefore more valuable than payout transaction data, although in reality the two complement one another to enable a host of richer analytical methods. As further chapters explore, gathering bet placement *and* payout data in the cryptocurrency domain, and indeed the blockchain games domain, is somewhat more complex.

4.4 Study 1: Systematic Review

The data that gambling activities generate can be used to build behavioural profiles of players in a number of ways. The most relevant to this thesis is the computation of behavioural measures which are derived from gambling

transaction data. In order to fully capture the range of behavioural measures employed across existing literature, a systematic literature review is required. This review should ideally capture the majority of studies found by existing systematic reviews, but should also capture studies published since publication of any existing reviews. The insights generated from this review will therefore provide a unique collection of commonly used computations. For example, authors may compute the sum of all bet sizes, known as the total amount wagered, as a proxy measure for gambling involvement [61], and base further analyses on that measure. It is also important to capture not only the way each measure is computed, but also the motivations behind each. These motivations, along with the broader contexts of the studies in which the measures are applied, will ultimately be used for framing comparisons between decentralised and centralised gambling in subsequent studies in this thesis.

4.4.1 Introduction

Systematic reviews are a widely recognised method of capturing the state of the art in an academic subject, and there are many ways in which they can be performed. At the highest level, a systematic review consists of utilising a scholarly search engine to gather all studies which match a given set of search terms. Once gathered, these studies can then be analysed, first to determine whether they meet some criteria of usefulness, and then their contents can be extracted. Several systematic reviews have already been conducted in the field of gambling studies, each aiming to describe some aspect of the field. For example, Chagas and Gomes (2017) review [1] critically analyses behavioural tracking research, outlining several trends in the methods employed, relationships between findings, and future directions of the field. They discuss some of the behavioural measures used in the studies identified, but do not focus explicitly on this aspect of each of the studies as they are instead more interested in the broader issues around gambling research such as self exclusion, behavioural feedback, and other wider issues. A more recent (2020) review by Lawn et al [7] takes a similar approach, outlining several more recent studies in the field but at a much higher level than specifically focusing on transaction analytics. Unlike Chagas and Gomes review, Lawn et al’s approach focuses on the identification of gaps in the literature, one such gap of note is that ‘the emergence of cryptocurrency

Table 4.1: Search terms used in Chagas and Gomes’ systematic review [1] (top) and the search terms used in the systematic review for this thesis (bottom). The key changes are the inclusion of the British spelling of ‘behavioural’, and ‘machine learning’ over ‘player card’ and ‘loyalty card’ to capture more technically oriented papers.

Study	Search Term
Chagas & Gomes	("actual" OR "behavioral* tracking" OR "tracking data" OR "big data" OR "real world" OR "player card" OR "loyalty card") AND "gambling"
This Study	("actual" OR "behavioral tracking" OR "tracking data" OR "big data" OR "real world" OR "machine learning") AND "gambling"

and block chain will likely also warrant research’. Studies of this nature are immensely important for developing a broader understanding of gambling research, yet in the context of cryptocurrency research where transaction data is the primary resource, a narrower and more specific focus is required.

Specifically, this study addresses RQ2 - Which behavioural measures can be applied in [new] domains. This requires understanding both the range of behavioural measures themselves, plus their relationship to any harm related variables such as self exclusion, as these relationships are what will ultimately inform interpretation of their application to the cryptocurrency domain.

4.4.2 Method

As Gusenbauer and Haddaway discuss, crawler based scholarly search engines such as Google Scholar are not best suited to perform systematic search tasks given their lower search precision and lack of features [62]. With this in mind, the Web Of Science search engine was chosen given its broad reach and good performance along their ‘necessary’ performance requirements [62]. The Web of Science functionality available for the search included the databases; Web of Science Core Collection, KCI-Korean Journal Database, MEDLINE, Russian Science Citation Index, and SciELO Citation Index. Unlike Chagas and Gomes’ review, only English language papers could be analysed, so no other languages were included. The Web of Science search tool also allows specification of document type, here only those listed as the ‘article’ type were taken forward. Finally, the publication date for the search was from 2007 (first publication in the field [61]) to 2021 (March 27th) inclusive.

The first step of this systematic review is to gather each of the studies

and identify whether or not their contents meet the minimum criteria of interest for this thesis. I use a modified variant of those used by Chagas and Gomes terms to adequately capture machine learning based studies since the publication of their review, given the development of the field outlined by Lawn et al [7]. I also applied a generalisation of the “behavioural tracking” term to include both American and British English spellings (See Table 4.1). With a collection of articles ready to be reviewed, studies pass the first stage if they use betting data to compute behavioural measures of some kind. Applying Fiedler’s observation that a behavioural measure is simply any variable derived from betting data [63], this first stage identification is fairly broad and still contains a number of papers whose contents will not ultimately be useful. This requires a first pass of reading to determine firstly if the study uses data or is simply a commentary or other type of article, and then whether or not the study computes any variables based on this data.

After the studies have been labelled, all relevant studies’ measures can be extracted. This step requires a thorough analysis of each of the ‘useful’ papers returned by the previous stage. At this point it’s important to note that the behavioural measures used between papers may be computationally identical but with different descriptors, or computationally different but with identical descriptors. Here, behavioural measures should be extracted based on how they are presented in each paper alone, i.e. the output of this stage is not a set of globally unique measures, but rather a set of the ‘raw’ representation of behavioural measures in existing literature. This raw collection can then be distilled by a process of repeated readings in order to identify which measures in the collection are in fact computationally identical. For example, Auer & Griffiths report the behavioural measure of ‘customer tenure’ [64], which is computationally identical (with a different unit) to that of ‘duration’ in earlier work by LaBrie et al [61]. At this final stage these two measures would be considered identical, and could be merged accordingly. The aim of this distillation is therefore to standardise the raw collection so that a complete picture of the development and use of different behavioural measures can be created. The stages of this review process are outlined in Figure 4.1.

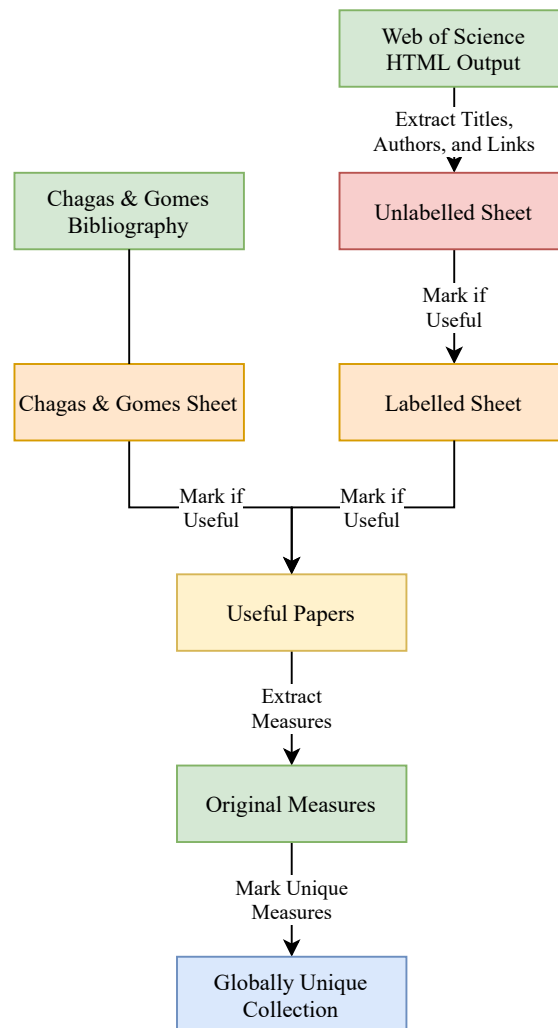


Figure 4.1: Each of the stages in the systematic review method, taking input collections of papers from Chagas & Gomes review, plus the up to date Web Of Science output. The result of this pipeline is a globally unique collection of behavioural measures used in gambling research.

4.4.3 Results

The Web of Science search using the terms presented in Table 4.1 yielded 469 results in total, 424 of which were classified as articles so were included in this review. Other document types which were not uncovered included reviews, meetings, and editorial material. Of the 424 articles found by the Web of Science search, only 47 met the minimum criteria described in Section 4.4.2. Additionally, of the 52 articles cited in Chagas and Gomes' original review, only 27 met this minimum requirement as many were commentaries or not transaction data driven studies. Given the similarity of the search terms and time scales in which the articles were gathered, it is not surprising that the 47 useful Web of Science articles plus the 27 Chagas and Gomes articles yield a total of 52 unique articles for further study in this review. This substantial overlap confirms that this more recent systematic review does indeed refresh the search process undertaken in Chagas and Gomes' work. With 52 unique studies identified between 2007 and 2021, the second stage of the review could be completed.

Removed Variables

The second stage of this review's methodology was to extract each of the behavioural measures used across these 52 unique studies. Given the broad scope of the first pass, not all of these studies used transaction data alone, and many used variables which were based on transaction data but were at the population group level, such as Ukhov et al's work exploring differences between sports and casino bettors [53]. Additionally, many studies such as those by Fiedler [65][63], Luquiens et al [66], and Auer and Griffiths [67], used game-type specific variables (See Section 4.2) such as the number of tables played, number of buy-ins, and other poker specific variables which cannot be generalised to other casino games.

In addition to these removals (group level variables and game specific variables), a further category of non-behavioural measure variables include those which exist at the account level, such as the number and size of deposits and withdrawals. These account level variables are important to developing a more holistic understanding of player's circumstances, but are not considered behavioural measures in the context of this thesis, so were removed. Studies which used account level variables included Haeusler's exploration of machine

learning methods to predict self exclusion [68], Braverman et al's study identifying high-risk gamblers [69], and indeed Ukhov et al's work mentioned previously [53].

Mixed Objectives

The collection of studies uncovered by this systematic review cover a broad range of objectives, and therefore employ a number of different experimental configurations. These varying study objectives include understanding behavioural differences between different types of gambling activities [53][70], measuring predictors of self-exclusion [71][72], broader studies around the impacts of personalised feedback on behaviour [73][74], and more. This mixed set of objectives is reflected in the broad range of behavioural measures used to investigate these issues, and is likely a driving force behind the use of functionally identical measures with varying textual descriptions.

In the context of this thesis, the most important aspect of each of the behavioural measures used is the motivation behind their use, or more plainly, why did the authors choose to compute any particular measure as a descriptor of gambling behaviour. The motivations behind each of the measures found will be more comprehensively discussed in the dedicated 'Behavioural Measures' section below (Section 5.1), but generally fall under aiming to either capture the financial or temporal involvement of the player with the gambling activity in question. For example, Auer and Griffiths note in their multi-publication discussion of their measure of theoretical loss [75] that their measure 'theoretical loss does not intend to cover other important aspects of gambling such as time involvement', also noting that 'gambling cannot be understood solely by looking at monetary aspects'. This view is generally held across the studies identified by this systematic review, with many studies computing both financially oriented and temporally oriented behavioural measures. Specific examples of relationships between different types of behavioural measure and conditions of interest, such as account closure due to gambling related problems, is discussed in Section 5.1 below.

Proprietary Collaborations

Several studies returned by the search conducted for this review included proprietary computations over player data which were not explicitly and

transparently defined. Studies which apply proprietary classification systems in peer reviewed research are not useful to furthering academic discourse as they cannot be replicated, so are not discussed in detail. These studies were Auer and Griffiths 2021 study on gambling before and after the global pandemic [76], the same author's 2020 work on personalised wager messages [77], and Challet-Bouju et al's 2020 work on generic player modelling [78], among others [79]. These studies use either the proprietary Mentor classification system², or the Playscan classification system³ [80], which are owned by Neccton GmbH and AB Svenska Spel respectively. It is however important to note that these systems each advertise a mix of survey and/or account level data, plus transaction data, in their computations. This mixed input means that even if their internal workings were known and could be replicated, their outputs cannot be considered behavioural measures in the context of this thesis, so are not of use anyway. A second important note from these publications is that behavioural measures derived from player transaction data are actually used in the computation of these classification systems [80], and that these classification systems are currently deployed across large populations of gambling operator customers to help identify players at risk of gambling related harms. This inadvertently supports a central tenet of this thesis; that measures derived from transaction data are in fact meaningful in identifying players at risk of gambling related harm. It does however also support the notion that it is important to incorporate multiple data sources in classification algorithms - a limitation discussed in detail in subsequent chapters.

Unique Measures

A total of 438 mixed-type variables were found across these 52 unique studies, 200 of which can be derived from a single player's transaction data. Of these 200 raw variables, 51 can be described as globally unique, and computed using betting transaction data alone, so are considered behavioural measures in the context of this thesis. Figure 4.2 charts each of these measures in relation to the data required to compute them, with the majority of measures derived from bet time data alone (See Section 4.3.2 above). This figure also visualises the interaction between the different data types discussed

²See <https://www.neccton.com/>, accessed 06/05/2021.

³See <http://playscan.com/>, accessed 06/05/2021.

in Section 4.3.2 above, and how they can be used to derive more complex behavioural measures than any single data type in isolation.

The 51 globally unique behavioural measures discovered by the iterative extraction process described above have not found equal application across each of the studies. Indeed, older behavioural measures have generally been applied more times than those created more recently, despite calls for more recent measures to replace their older counterparts [75]. This is to be expected and caveats the interpretation of the number of uses of each measure as a proxy of their effectiveness or reliability. The distribution of the number of uses each of these measures in the studies collected is available in Figure 4.3. Given the importance to this thesis of the collection of measures uncovered by this review, a complete discussion of each of the behavioural measures is deferred to its own section (5.1) in this thesis. The following discussion is therefore specific to the execution of this systematic review alone.

4.4.4 Discussion

In the context of the size of the gambling industry globally, the number of studies which use actual gambling transaction data to develop any understanding of players in any capacity is somewhat underwhelming. In addition to this lack of quantity, many studies utilise data sets which are not publicly available, therefore even the published and peer reviewed results are impossible to verify. This lack of quantity and transparency directly affect the number of behavioural measures which can be applied to new domains, such as those later in this thesis, but may be as much an artefact of the relative youth of this field as it is a reflection of the transparency and collaborative drive of the industry.

This considered, the set of 438 non-unique variables found in the 52 studies does indicate a somewhat rapid pace of innovation within the field itself, with researchers quickly branching out to uncover new relationships and descriptive methods for players in general. The 51 globally unique behavioural measures in particular reflect this innovation, although as the graph (Figure 4.3) of their distribution within the literature shows, many are not yet mature enough to be meaningfully applied for comparative purposes in new domains. For example, the behavioural measure of ‘mean amount wagered per session’ has only been used once (that this review returns), which

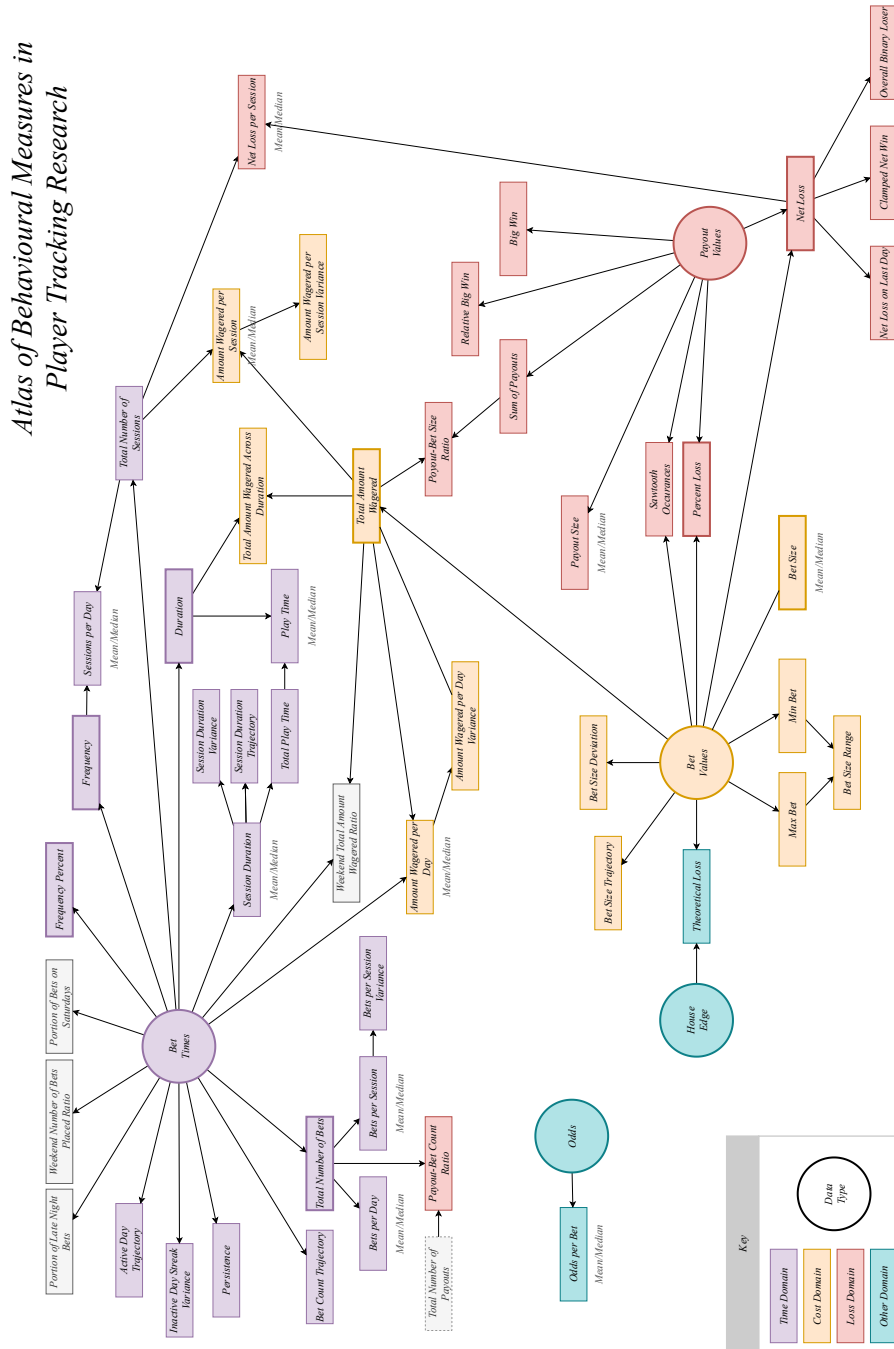


Figure 4.2: Globally unique collection of behavioural measures identified in Study I in this thesis. Colour coding indicates the domain in which each behavioural measure resides, this is based on the information needed to compute each measure - a concept expanded upon below in Section 5.1. Greyed out measures required contextual information (time zone), and arrows signify a computational dependency between measures, e.g. Total number of sessions is required to compute amount wagered per session.




Figure 4.3: The distribution of behavioural measures used in studies uncovered by the systematic review. It is important to note that the raw number of times a behavioural measure is used in the literature is not necessarily and indicator of its validity, but acts as a useful starting point for further investigation.




Figure 4.4: Each of the unique behavioural measures uncovered by the systematic review, plotted against the studies they were used in. Note that in some instances in text a measure may be referenced as being used in two studies but appears as a three in this figure. This is due to variants or transformations of the same measure being used in a single study.

means that if one were to compute this measure in a new domain, it would be difficult to compare the two values in any meaningful way. Naturally this does not discount the approach one may take of applying all of the behavioural measures to the new domain under the rationale that they may one day be meaningful for comparison, but this exploratory ‘apply everything’ approach may not make the best thesis.

Limitations

This systematic literature review, whilst gathering a large sample of studies and behavioural measures, has several important limitations. The most significant limitation is that only one scholarly search engine (Web of Science) was used. While this search engine does cover a number of different databases, incorporating multiple search engines would naturally return more articles at the obvious cost of additional search and analysis time. An alternative to increasing the number of search engines used would be to take a publication-oriented approach based on an existing systematic review. For example, one may take the venue of publication of each of Chagas and Gomes’ cited studies and search each of them for new and useful articles. This approach may however miss work published farther afield, and in newly created venues, so was not used for this review.

An additional limitation of this systematic review is that although the term ‘machine learning’ was included in the search terms, many of the now widely recognised machine learning based studies in the field of gambling research were not successfully returned. Examples of work that were not returned by the search include the seminal study on data mining techniques and player data by Philander [81] and work by Sarkar et al which applies knowledge extraction techniques to machine learning methods for safer gambling [82]. This non-exhaustive set includes studies which have been extremely important to the development of machine learning applications in gambling research, but which were not explored as part of this review.

One reason that these studies were not returned by the search is that they may use more specific terms than simply ‘machine learning’ in their titles and keywords. The remedy for this would be to further expand the search query to include all specific terms of interest, however this broadening of search terms - especially in the direction of machine learning research - would lead to the return of a prohibitively large number of papers, particularly

across unrelated fields such as machine learning applications for AI in digital games, for example. A second reason several machine learning oriented studies may not have been returned is that they simply don't exist in the databases indexed by this particular scholarly search engine so would not have been returned no matter the query design. In any case, a set of known (and missed) machine learning oriented studies will be referenced throughout the dissection of behavioural measures below for completeness. These have been uncovered by manually searching for citations of Philander's seminal work on data mining [81].

Outcome Variables

A final limitation of this review, and indeed of the studies identified across the field of player behaviour tracking, is that the choice of outcome variables is not consistent, and that of the outcome variables chosen, some may lack broader validity when used in different contexts. Specifically, several studies use outcome variables such as likelihood of self-exclusion or self-reported problem gambling severity, which naturally rely on the affected individual being aware of their circumstance and take proactive steps to mitigate gambling's effect on their lives. Using such variables as outcomes in a study may mean missing the sub-population of players who experience some kind of gambling related harm but are unable to self-identify (and subsequently seek help) and therefore continue to experience harm. It is by identifying not only those who experience harm and act on it, but all those who experience harm, that a more holistic understanding of gambling behaviours can be created. This is not to say that these studies are not useful in the context of this thesis, but rather that they contribute to a subset of our understanding of behaviours and should be considered in the broader framing of outcome variable validity.

4.5 Summary

This chapter has discussed types of gambling, types of game, and the types of data they generate, which serve as essential background for understanding a field of research within gambling studies known as player behaviour tracking. Player behaviour tracking concerns understanding players using their transaction data. This can be achieved by computing behavioural measures

across player's betting data, and then applying statistical analyses to these measures. In order to fully understand the range of different measures used across this research, a systematic review was conducted, finding a large collection of studies which employed a range of different measures. The entire next chapter is dedicated to exploring the range of behavioural measures uncovered by this systematic review, and delves deeper into the motivations behind each of these measures and their implementations.

Chapter 5

Understanding Behavioural Measures

“Excuse Me.”

Darth Jar Jar Binks
The Phantom Menace

The previous chapter introduced the field of player behaviour tracking, including a discussion of the study-level results of a systematic literature review. This chapter delves deeper into each of the studies returned by the systematic review, focusing explicitly on which transaction level behavioural measures are used, and the motivations behind their use. It begins by describing the broad differences within existing research in Section 5.1, such as the different scales in time used, parametric and non-parametric equivalents, and inter-period derivatives of the same measures. These broad differences are important to discuss here as they affect how these behavioural measures are interpreted when applied. It then describes each of the behavioural measures discovered in Section 5.2, dividing the collection of measures by the information required to compute them. This information-oriented taxonomy was chosen over alternatives (such as type of gambling activity studied) as it maps directly to the process of decoding cryptocurrency transactions as described in Chapter 3. For example, Section 5.2.1 describes the time domain, whose behavioural measures require the time of each bet placement. This means that by retrieving the time of each bet placement transaction, all of the measures in the time domain can be computed.

5.1 Existing Research

Study 1 above uncovered a collection of papers, each applying some kind of computation to a set of transaction data. These computations typically aim to derive a single value from a set of data, commonly referred to as a behavioural measure, and relate that measure to some internal or external condition of interest such as self exclusion, or self reported problem gambling severity, for example. These behavioural measures therefore contribute to developing a behavioural profile of the players in each of these studies, inviting the question of exactly how many unique behavioural measures have been used, how they relate to one another, and how they relate to the external conditions of interest.

This section expands the findings of the systematic review by decomposing the collection of behavioural measures into groups by the level of information required about each bet required to compute each measure. This implementation-oriented decomposition is used over alternatives, such as grouping by the type of gambling activity the measures were used on, in order to develop a more practical understanding of exactly which data is required - as in the context of this thesis this may become a limiting factor. Motivations behind each of the behavioural measures are also discussed, including any work exploring their correlation with other measures where possible. Providing a broader motivational backdrop for each of the behavioural measures allows not only a deeper understanding of how they are conceptually similar, but is key to interpreting their results when they are computed across a set of actual transaction data in a new domain in the studies below.

Additionally, it is important to note that although the following discussion outlines each of the behavioural measures in isolation, and their contribution to external harm related variables, it is the emergent profile generated by the use of multiple measures at the same time which provides the richest descriptions of player behaviours. The reductionist approach of answering which measure is most important or descriptive for capturing a given aspect of player behaviour is then less meaningful. Instead, uncovering which collection of measures accurately and broadly relates to some external variable in the most reliable way helps build a more holistic picture of the different types of behaviours that can exist in gamblers. To this end, each of the descriptions of the behavioural measures below also contains a brief discussion of its

relation to other measures where known.

5.1.1 Study Variation

Although the collection of measures themselves can be broken down by the information required to compute them, several notes can be made about the variation in the studies themselves and how these variations affect the interpretation of the behavioural measures in question. Specifically, the set of studies returned by the systematic review varied in terms of the scales in time that they used, the use of parametric vs non-parametric statistics, and the use of inter-period derivatives (which are simply measures computed across different time periods for comparison). This subsection briefly discusses these three aspects.

Scales in Time

The issue of the many different scales in time that studies can explore must be addressed, as this has a significant effect on the way behavioural measures are interpreted (e.g. frequency per day or frequency per month), and the types of measures that can be computed. At the highest level, the scales of time under inspection in gambling research vary from the individual session level, to the weekly or monthly scales, up to describing and comparing entire gambling careers. As discussed briefly above, this can become entangled with the types of measures computed, since a study may be exploring gambling behaviour across an observation period - or scale - of two years, but compute behavioural measures such as mean session length. Conversely, a study may focus on understanding within-session behaviour, but may also include variables at the career level. In order to avoid any confusion in the following discussion, and across each of the studies presented in this thesis, the following naming convention is used; **Time Scale** refers to the total period of data available for a given study. **Time Domain** refers to the group of behavioural measures which may be computed using bet-time based data - e.g. frequency in the 2 week time domain.

It is also common to encounter studies which are concerned with a single time scale (e.g. gambling involvement over a period of 6 months), but which apply behavioural measures from multiple time domains (e.g. session level variables, monthly variables, and career variables). Possible scales in time

include the individual bet level (e.g. [83]), any arbitrary length of time (e.g. one week [84], two weeks [85], three months [86], etc), empirically determined lengths of time such as sessions [65][63], or entire player careers [87][61]. This rich range of study time scales presents many opportunities for verifying or challenging findings in across the different scales, but in the context of discussing the computations and motivations behind the behavioural measures themselves is not discussed further. This considered, the scales in time in which studies exist is vital when framing comparisons in their findings with similar studies in new domains. The studies presented herein will explicitly address these varying scales when comparisons are made, and interpret results accordingly.

Parametric vs Non-Parametric

Many of the measures encountered in the systematic review can be described as parametric or non-parametric equivalents of the same fundamental computation. For example, the behavioural measure of session duration can be described by its mean or median value, since the computation of session duration for a given player is itself a sequence of values. One important observation is that mean values can be derived from aggregate sequences of bets, but median values require knowledge of the individual transaction data. For example, the mean bet size can be computed as the sum of all bet sizes divided by the number of bets placed, whereas the median bet size cannot. In data sets where these values are already aggregated and the individual bets are not available, such as many of the *bwin* data sets used in the Harvard studies [61][2][88], this means that median bet sizes cannot be computed. These types of parametric and non-parametric behavioural measures are treated equally in this chapter and indeed this thesis, but it should be noted that actual gambling transaction data is rarely normally distributed, so non-parametric measures are typically preferred. This preference is well documented [88], and motivates the use for non-parametric statistics throughout the work presented in this thesis.

Inter-Period Derivatives

Finally, inter-period derivatives can be described as the computation of the same behavioural measure across different time periods, with the ultimate

goal of deriving some metric of change between the two. For example, Percy et al's work predicting online gambling self exclusion [72] computed five key behavioural measures. They also applied several transformations to these measures to create derivative behavioural measures which described changes between periods of time. An example from their work is the measure of bets per day, which meets the definition for inclusion in this thesis described above. However, further derivatives of this measure such as the statistical significance of the change in this measure between two periods of time, as used in their paper, sits one level higher in the analytical sense than the behavioural measure itself as it is contextually dependent on the experimental configuration of their study. These types of inter-period derivatives have been removed during the data cleaning process described in Section 4.4.2 as they depend on external information (experimental configurations) in addition to the raw transaction data.

5.2 Behavioural Domains

Following extraction of each of the behavioural measures from the systematic review described in the previous chapter, it became apparent that the set of measures used across player tracking research naturally fall into a groups along two axes. The first axis is the scale in time in which the behavioural measure exists. For example, total amount wagered uses the exists in a time scale equal to the entire (continuous) duration of the study, whereas amount wagered per week naturally exists at a (discrete) weekly scale. The second axis is the information required to compute each of the behavioural measures. The following sections break down the behavioural measures identified along the information required to compute them, then by the scale of time they exist within.

5.2.1 Time Domain

Behavioural measures in the time domain aim to capture temporal patterns of play in some way, and are typically motivated by a desire to operationalise the temporal aspect - as opposed to the financial aspect - of gambling involvement [89]. Measures therefore exist in the time domain if they can be computed using only knowledge of the times that bets have been placed. Formally, behavioural measures in the time domain can be defined as any

value derived from the finite sequence T , where $T = (t_0, t_1, t_2, \dots, t_n)$. Here T is simply a sequence of timestamps. Following the discussion of transaction level data and the anatomy of a bet in Section 4.3.2 above, this is one of the core components of betting transaction data, and can be applied to both betting and payout transactions. Importantly, this means that behavioural measures based on bet times can be applied to payout times to generate similar but mirrored variables, although in practice bets and payouts are typically treated as a single transaction for analytical purposes. Measures in this domain also include all of those which can be computed using any derivative of T , for example, each of the times t_x can be grouped into sessions, or into discrete bins equating to a fixed time period e.g. days, weeks, and so on. Each of these derivatives of T bring their own limitations, so have their own dedicated sections below.

With this in mind, and in order to capture these patterns more accurately, measures in the time domain can be split into the continuous, discrete, and session sub-domains. Continuous time domain measures view the entire collection of betting data as a single collection, or formally, use the values in the sequence T . Discretising this sequence means grouping bets into bins corresponding to fixed time intervals, thereby creating a sequence of sequences. Behavioural measures in the discrete time domain therefore require transaction level data at a minimum, or in cases where measures are based on sums of these discretised bins (such as total amount wagered per day) require the aggregate values of these discrete bins. Finally, grouping bets by temporal proximity with some fixed time window yields the session subdomain. The following subsections outline each of these sub-domains in detail, and the measures uncovered by the systematic review which fall into these domains.

Continuous Time Domain

Behavioural measures computed in the continuous time domain treat the whole sequence T as a single and complete collection. This highest possible level grouping allows the computation of measures across the maximum possible length of observation for each player, and therefore holds commonly used behavioural measures uncovered by the systematic review. The very general nature of this domain does however mean that there are a finite number of computations that can be performed, and therefore a small number

(2) of behavioural measures in this specific subdomain.

The most commonly occurring behavioural measure in the continuous time domain is that of **duration**, which is simply the difference between the first and last elements of T , or $t_n - t_0$. This makes duration a useful descriptor at the aggregate level when establishing differences in observational periods between participants for data gathered over a fixed period of time. Given its career-level nature, it does not capture any granular details about the player's interaction with a gambling activity, but has seen widespread use, especially in studies using the Harvard-affiliated bwin datasets [2][3][88], with its first use by LaBrie et al [61] in 2007. Including the Harvard series of studies, duration has been applied in a total of 16 studies across a number of different types of gambling. Duration is an important behavioural measure for providing temporal context to more granular measures, and can also be used to naively distinguish between levels of temporal involvement in data sets over a fixed time period. As Fiedler notes in his longitudinal studies of Poker players [65][63], many of the behavioural measures that appear below can be combined with duration to provide richer descriptions of player behaviours.

The **total number of bets**, computed as simply the length of sequence T , is another common behavioural measure which captures the level of gambling involvement per player in a very simplistic way. As with the behavioural measure of duration, the total number of bets was first used by LaBrie et al [61], and has been used 18 times in the studies returned by this review. The total number of bets has been positively related to structural game characteristics (rewards and betting options) by Leino et al [90], making it most meaningful when used to compare between two similar types of game. Additionally, the total number of bets has been found to be substantially higher in so-called 'heavily involved bettors' [61][91], which makes intuitive sense given 'heavy involvement' is typically a top percentage of players by total amount wagered. This relationship - between total number of bets and total amount wagered - is discussed below.

The career level nature of analysing a player's betting activity as a single continuous sequence has several limitations. The most important of these limitations to note is that any measure at this level by its design holds little descriptive power over more granular time scales. For example, the measure of duration gives no insight into the actual play time of a given player, and

the total number of bets gives no insight into the frequency of betting. For more granular insights, static discretisation (fixed time intervals) or dynamic discretisation (partitioning into sessions) is required.

Discrete Time Domain

Discretising a continuous series such as transaction data invites the development of more nuanced behavioural measures at the cost of granularity, as the discretisation process naively merges contiguous observations. As mentioned above, there are several levels of discretisation, the most common of which being daily bets. This aggregation results in one sequence D which contains bets for each calendar day upon which at least one bet was placed. For example, a player placing six bets split equally across two separate calendar days would yield a sequence D of length 2, where each element of D would itself be a sequence containing three timestamps. Equations 5.1 and 5.2 show how the sequence of six bets can be split this way.

$$T = (t_0, t_1, t_2, t_3, t_4, t_5) \quad (5.1)$$

$$D = ((t_0, t_1, t_2), (t_3, t_4, t_5)) \quad (5.2)$$

The most commonly used behavioural measure in the discrete time domain is that of **frequency**, which can be described as the length of the derivative sequence D . In other words, this is the number of calendar days upon which at least one bet was placed. This measure appears in over 20 of the studies uncovered by the systematic review, making it one of the top 3 most commonly used in the field. First used by Xuan and Shaffer [92] as a purely descriptive variable, frequency has since been used by Dragičević et al [93] as a ‘risk factor’ in their cluster analysis of casino game players, finding that two clusters ($n=80, n=6$) exist in their sample ($n=546$) which exhibit extremely high values along this measure. This matched Braverman and Shaffer’s earlier study using an identical methodology [3]. These early studies suggested that frequency - an obvious metric of time involvement - could be used to make meaningful distinctions between groups of players. This idea matched work by LaBrie and Shaffer in 2011 [94] which found that betting frequency was higher in players who closed their accounts due to gambling related problems (as opposed to other reasons). This finding was confirmed

by Gray et al [95] in the context of account closure vs non-closure, although is challenged in Gray et al's later work [91] which divided players into groups based on total amount wagered and total number of bets. These findings taken together suggest that the naive approach of taking the top percent by any single metric may not be meaningful in the context of understanding the potential for gambling related harms, but instead functions as a purely descriptive analytical technique.

The measure of **frequency percent** is an example of combining an existing behavioural measure with that of duration in order to derive a normalised value across players. Formally, frequency percent is simply frequency above represented as a percentage of duration. Although frequency percent appears to simply be an alternative representation of frequency, the use of duration in its computation makes it a unique behavioural measure in its own right in the context of this thesis. Frequency percent has been described using a number of different terms such as 'frequency of betting' [96] in the studies returned by the systematic review (and is often conflated with just frequency). Frequency percent is less common than frequency, with a total of 13 uses overall, starting again with LaBrie et al's 2007 study [61] through to Edson et al's 2021 study [97]. Nelson et al's 2008 study [98] found that frequency percent was significantly different (higher) in self-limiting players (pre limit setting) than non-self-limiting players (players who choose to add financial limits to their accounts) across sports betting participants. Similarly, Gray et al [95] in addition to computing frequency also computed frequency percent, finding it was also slightly greater in account closure cases vs non-closures. These studies, and others returned by this review, provide empirical support for Currie et al's assertion that 'the chances of experiencing gambling related harm increase[d] steadily the more often one gambles...' [8]. This makes intuitive sense, and shows that both frequency and frequency percent can be meaningfully applied to describe behaviours in new domains.

Closely related to the concept of duration is that of **persistence**, which is defined as the number of consecutive months with at least one active day (one bet). This more nuanced variant of duration naturally places emphasis on continuous play, so should in theory be more descriptive than measures like duration when comparing external variables such as self exclusion. Unfortunately, very little work has explored this behavioural measure, as the

first and only paper to use is was Edson et al in 2021 [97]. It is therefore unclear how persistence relates to gambling related harms, and any proxy measures thereof such as self exclusion, ultimately making it difficult to apply to new domains.

Similarly, **Active day trajectory** is one of a family of measures which compute the trajectory of the values in a sequence. This means using the slope of a linear regression fit to the sequence as the measure, with active day trajectory referring to the slope coefficient fit to the number of active days across time periods. For example, Braverman et al use this measure in their 2013 study on early identification of high-risk internet gamblers [69]. Their study splits a player's wagers in half, computing the trajectory between the two halves and categorising the result as one either increasing, decreasing, or stable. Unfortunately, as with other behavioural measures which have only been used in a single study, it is difficult to concretely map the distribution or interaction of this measure to a harm related external variable. This considered, Braverman et al did include the measure in their final classification tree (Chi-square Automatic Interaction Detection tree¹), which in context means that it has meaningful power in predicting whether or not a user will be flagged by the *bwin.party* responsible gambling program as having gambling related problems. Nonetheless, the strength of this power is unknown, and has not been confirmed in any similar settings by any subsequent studies. This means that the measure of active day trajectory alone cannot be used to provide meaningful distinctions between groups of players in new domains, but this does not mean that it is not still informative.

The measure of **inactive day streak variance** refers to the variation in days between wagers. As with many other measures uncovered in this review, this measure has only been applied once - in a recent study by Ukhov et al [53] exploring differences between casino and sports bettors. Their study reveals that the inactive day streak variance is only slightly meaningful in explaining the risk of exclusion due to problem gambling. This weak relationship (Median normalised absolute Shapley value of 0.2 for casino players and 0.05 for sports bettors) extracted from a classification model (XGBoost by Chen and Guestrin [99]), means this behavioural measure is generally of little predictive power, and is not of use in the context of this

¹This is one of the oldest classification tree algorithms available, and has not been used in any of the other studies returned by the systematic review.

thesis.

The behavioural measure of **bets per day** is one of the most commonly used measures in gambling literature, and is simply the mean or median of the lengths of each sequence of bets per day in D . Intuitively, bets per day captures the amount of gambling occurring, albeit in a more mechanically focused way than measures like frequency above. With a total of 19 uses starting with LaBrie et al's seminal work [61], high values of this measure have been positively related to responsible gambling events [95], presence in heavily involved subgroups in casino players [2], and likelihood of being flagged by the *bwin.party* responsible gambling program as having gambling related problems [69] to name a few. As with frequency and frequency percent above, this makes intuitive sense given the nature of gambling related harms, and makes bets per day a strong candidate for application to new domains for exploring the scale of potential harms.

The **bet deviation per day** refers to the standard deviation of the number of bets placed across all of the active betting days in a transaction sequence. It therefore partially captures both the financial and temporal aspects of gambling involvement, and was alternatively described as 'bet variability' when it was used by Braverman et al in 2013 [69]. The first study to use this behavioural measure returned by this systematic review was Dragičević, Tsogas, and Kudic's 2011 study on behavioural risk markers for high-risk gambling [93], although the first ever study to use this measure was Braverman et al in 2010 [3]. Having only seen use twice, bet deviation per day has not been applied broadly enough to reliably relate it to factors such as responsible gambling events, however it has been used in the context of behavioural clustering - a type of study described in a dedicated section (5.3) below.

The second trajectory oriented measure is that of **bet count trajectory**, which requires fitting a linear regression to a sequence of bet counts aggregated to discrete intervals. The measure, like the active day trajectory described above, is the coefficient of this model, although this value can also be categorised into one of 'increasing', 'decreasing', or 'stable' depending on its value (positive, negative, zero, respectively). As with the previous measure, bet count trajectory was used in Braverman et al's 2013 study [69], although it is important to note that this is not the same as bet size trajectory - a similar but financially-oriented measure in the bet domain which has been

used in similar work. Unlike the other measures used across Braverman et al's studies, bet count trajectory has not been used in any other studies returned by this systematic review, which unfortunately means that little information exists around its relationship to other behavioural measures, and indeed to other external variables of interest such as potential for gambling related harms. This does not mean that these insights are not informative, however. For example, the 'trajectory of bets' as they describe it is included in their final model for the early identification of high-risk internet gamblers for fixed odds betting games, making it a potentially fruitful area of further investigation in new labelled data sets. That considered, the lack of studies which use it means it is ultimately not useful in further analyses in this thesis.

Session Domain

Behavioural measures relying on session level data require partitioning the continuous time domain into chunks, where each chunk contains a set of bets in close temporal proximity to one another. This proximity, typically determined by a fixed time value, is referred to as the session window. Session windows in existing work range from 10 minutes [66] to 30 minutes [65], the exact value to use is a decision is typically made at the discretion of the authors rather than by any empirical means. This variable session window can ultimately impact the distribution of any session level behavioural measures, an issue highlighted as encountered throughout this thesis.

The most obvious behavioural measure in the session domain is that of the **total number of sessions**. First used by LaPlante et al in their study of poker players [100], this measure is formally defined as the length of the sequence S . Since its conception it has been used a total of 5 times, in subsequent studies of poker players [65] [63], in Ukhov et al's work comparing behavioural profiles of casino players and sports bettors [53], and in Finkenwirth et al's study of online gambler's self exclusion [101]. Of the five studies which have employed the use of the total number of sessions as a behavioural measure, Finkenwirth et al found that it had little importance as a feature for predicting voluntary self exclusion in a balanced sample of 2,157 self excluders and 2,157 non self excluders [101]. Similarly, Ukhov et al found it 'less informative for sports betting compared to casino gambling', in the context of its relationship to a risk score developed to operationalise

likelihood of self exclusion related to problem gambling. These findings likely stem from the high correlation between duration and total number of sessions discovered in poker players by Fiedler [63]. This makes intuitive sense, as the longer a player is gambling over time, the more sessions they are likely to play. Subsequently, given the aggregate nature of the measure, it is unlikely to be useful as a distinguishing factor between self exclusion and non self exclusion as it is unable to capture more nuanced differences in player profiles. This means its application to new domains may not provide any meaningful insight which simply using duration would reveal.

The first behavioural measure to make use of both the discretised time domain and the session domain is that of **sessions per day**. Formally, this is the mean or median of the number of sessions played per active day, which is the elements of sequence D themselves decomposed into elements s representing each of the sessions within that day. First used by LaPlante et al in their 2009 study of internet poker gambling behaviour [100], the behavioural measure of sessions per day has since been used a handful of times, especially in studies focusing on poker transaction data [65][63]. Despite it's poker-centric origins, the measure of sessions per day has been transferable to gambling more generally, as Dragicevic et al showed in their 2015 research into self exclusion [102], and again in Finkenwirth et al's 2020 research [101]. Despite these more recent and non-poker focused studies, little information about sessions per day exists in the context of simple casino games, making it difficult to find meaningful application in the context of this thesis. For this reason sessions per day is not explored further.

The measure of **session duration**, also known as play time, is the mean or median value of the lengths of each of the sessions a player has had. Formally, this a measure of centrality of the sequence of lengths of each session ($t_n - t_0$ for each t , for each session). It therefore aims to capture the typical time commitment of a player to a gambling application, although in practice only the parametric variant of this measure has been used, which can easily be skewed by extreme values. (mean) Session duration was first used by Fiedler in his 2011 study of poker players [65], and has been used a total of 3 other times in the papers returned by this systematic review. These include Fiedler's later poker study [63], Percy et al's exploration of predictive machine learning models [72], and Finkenwirth et al's study of self exclusion [101]. Like other session domain measures, session duration

has simply not been applied broadly enough to make its application to new casino game domains meaningful, so again is not explored further.

A complementary measure following the typical time commitment of a player above, is how varied this time commitment is across all of their gambling sessions. The measure of **session duration variance** captures exactly that by computing the variance (in the statistical sense) of the sequence of session durations, computed as described in the previous measure. Unlike other measures in the session domain which have been used in analysing poker data, session duration variance has only been used twice in the studies returned by this systematic review - first by Ukhov et al in 2020 [53], and then in the same year by Finkenwirth et al [101]. Both studies take a predictive approach to using behavioural tracking data, using a battery of behavioural measures and performing analyses across them. Ukhov et al find session duration variance to be only weakly useful (#12/40) for explaining gambling related exclusion in casino game players, but find it to be the 6th most powerful variable (in their set of 40) for their sample of sports bettors. Finkenwirth et al however, apply the measure of session duration variance but do not discuss it in their paper, making it difficult to conclusively determine its relation to any external harm related variable.

Like other trajectory type measures, the **session duration trajectory** is the coefficient of a linear regression, but here it is on the lengths of each session. This measure can therefore be most intuitively described as a players change in temporal gambling involvement across sessions, but has a weakness it shares with all of those in the session domain that a player may go a long time between playing, thereby potentially skewing this value. This measure has been used by Ukhov et al [53] and Percy et al [72]. Ukhov et al find that the session duration trajectory is important for predicting problem gambling in casino players, although as discussed above their model is somewhat poorly fit to their data. Percy et al instead found that there was no substantial difference between a self excluding cohort and a control group, and that the self excluding cohort actually averaged (mean) a value of 0.0 for this measure. These results are therefore somewhat opposed, drawing into question the meaningfulness of this measure's application in new domains.

At the more granular scale in the session time domain are the measures of **bets per session**, and **variance of bets per session**, which have been used a total of 2 and 1 times respectively. As the names suggest, these

each capture the centrality and variance of the number of bets per session. Both are applied by Finkenwirth et al [101], with bets per session being additionally applied by Sagoe et al [70]. Finkenwirth et al found both to be less descriptive than alternatives in their sample, suggesting they explain less than 3% of the variance between voluntary self excluders and non-self excluders. Sagoe et al focussed on behavioural effects of physical gambling venues, so neither measures nor this particular study are discussed further.

The final behavioural measure in the session time domain is that of **total play time**, which is the sum of session durations. With a total of 6 uses it has been gaining adoption following its introduction by Fiedler in 2011[65], including by Auer & Griffiths study on limit setting [75], Dragicevic et al's work on self exclusion [102], Sagoe et al's work described above [70], Ukhov et al [53], and a second time by Fiedler [63]. Dragicevic et al found only a small difference (mean 19.5 versus 19 hours, median 6.0 versus 5.5 hours) in total play time between self excluders and their control group. Both Fiedler, Sagoe et al, and Auer & Griffiths's studies used total play time in a descriptive context rather than in relation to a harm related variable, hence it is difficult to support or refute Dragicevic et al's finding, and therefore meaningfully apply it to new domains. For this reason, we can now move to the second major behavioural measure domain; the bet domain.

5.2.2 Bet Domain

In addition to measures identified in the time domain, measures in the bet domain require information regarding the size of each bet placed. Following the sequence notation used above, this gives sequence $B = b_0, b_1, \dots, b_n$, where b_x is the size of a bet in whichever currency is being used as stake. As in the time domain, the bet domain can be treated as a single continuous collection, discretised using a fixed time window, or dynamically discretised into sessions using a session window. This subsection presents each of the bet sub-domains and their behavioural measures.

Continuous Bet Domain

The most commonly used behavioural measure is that of the **total amount wagered**, which is operationalised as the sum of the sequence B . Having been used over 30 times across a range of different studies, the behavioural

measure of total amount wagered is the most intuitive measure of financial involvement in gambling and has been positively linked to a number of harm related variables. First used in LaBrie et al’s seminal work [61], total amount wagered has also been used as a population grouping measure in its own right. For example, epidemiological studies by LaBrie et al [2] and LaPlante et al [100] have used top portions of players along this metric as distinct ‘heavily involved’ groups. As discussed in Subsection 5.2.1 above, this naive grouping may not be informative for understanding potential harms in the population, but does provide an important broader context for understanding how the population is composed.

When used to explore links between total amount wagered and harm related variables, Gray and Shaffer found this measure to be dramatically higher in players who trigger responsible gambling interventions than in their control group across fixed-odds sports betting, live-action sports betting, and casino betting [95]. Similarly, Catania & Griffiths found (mean) total amount wagered across a their sample of voluntary self excluders to again be dramatically higher in those who self-reported gambling addiction versus other self excluders [71]. Conversely, Finkenwirth et al found a variant of this measure - total amount wagered from promotional offers - to be the fourth most important feature in their random forest model for predicting self excluders, although total amount wagered itself was not found to be useful [101]. The findings that total amount wagered or a contextual derivative thereof appear to be related to a number of harm related variables makes intuitive sense. As described when discussing frequency percent above, this supports the second part of Currie et al’s assertion that ‘the chances of experiencing gambling related harm increase[s] steadily ... the more money one invests in gambling’[8]. The behavioural measure of total amount wagered is therefore considered essential to understanding potential harms in new domains, and will be used where possible in later studies in this thesis.

The behavioural measures of **bet size** and **bet size deviation** capture both the central tendency and spread of the values in the sequence B , and have seen use 14 times and 5 times respectively. As the 7th most commonly used behavioural measure, (mean) bet size has been used both to provide epidemiological descriptions of populations of players [2][100][61][103][91], and in relation to harm related variables such as self-limiting behaviours [98][104], account closure due to gambling related problems [105], and trig-

gering corporate responsible gambling interventions [95]. Nelson et al found that (mean) bet size alone is not meaningful as a predictor of self-limiting behaviour[98]. Although Broda et al found that (mean) bet size was higher in both fixed-odds and live-action bettors who exceeded their deposit limits. This finding is supported by Gray et al’s study on corporate interventions, who found that both mean and median bet size was also generally higher in positive responsible gambling cases [95]. Braverman and Shaffer [105] unfortunately don’t discuss (mean) bet size in their results despite describing it in their methods. From these studies it is clear that bet size may not be the perfect behavioural measure for identifying the potential for gambling related harms, but it appears to be informative nonetheless so is likely useful in the context of understanding new domains.

Bet size deviation has seen less use, first used by Braverman and Shaffer [3], then by Dragicevic et al [93], Adami et al [106], again by Braverman et al [69], and finally by Percy et al [72]. These studies each use focus on the use of unsupervised machine learning to address gambling specific research questions, finding that bet size deviation (also referred to as variability) is an important measure in the identification of self-reported gambling related problems [3]. Dragicevic et al’s findings largely echoed those of Braverman and Shaffer, identifying a single behavioural group within their sample who exhibited extreme values along this measure [93]. In Dragicevic et al’s study, the group scoring highest in this measure also exhibited the largest mean net losses of €21,650 in comparison to €2,570 over a one year period. Similarly, Adami et al’s study into the behavioural clustering of problem gamblers applies bet size deviation amongst a number of other measures, finding that it is substantially higher in one of their clusters, which contained 9 problem gambling cases out of 16 members. This was the highest problem gambling concentration proportionally (56%) of all groups in their study. These studies all add to the notion that not only is bet size deviation important to the identification of potential harms using transaction data, but that behavioural groups exhibiting high values in this measure may be at a greater risk than the general population. Bet size deviation is therefore an important measure to compute as a harm related proxy, and less as a generic descriptive variable such as duration.

With the centrality and spread of the sequence B captured, another derivative is that of **bet size trajectory**, the fourth ‘trajectory type’ measure

uncovered so far. The earliest use of bet size trajectory was by Braverman and Shaffer in 2010 [3] in their study of behavioural markers for high-risk internet gambling. In their work, the bet trajectory is applied at the aggregation level of daily bets and at the one month time scale. This level of aggregation was a limitation incurred by their data set which only included daily aggregate data, hence they were unable to determine if trajectory changes were caused by an identical number of larger bets or by more numerous bets of the same size. Nevertheless they find that bet size trajectory was highest in a high activity, high variability behavioural cluster of which 73% of members reported closing their account due to gambling related problems. Dragicevic et al's subsequent study found mixed results regarding bet size trajectory, with a high frequency behavioural cluster exhibiting only moderately positive values for this measure [93]. Similarly, Adami et al's application of bet size trajectory appears to be inconclusive for distinguishing between distinct behavioural clusters including problem gamblers [106]. Percy et al on the other hand found bet size trajectory to be notably different between a self excluding cohort and control group (19 versus -50), although this finding was not statistically significant ($p=0.2$). Bet size trajectory therefore cannot be conclusively related to any harm related variable, but has seen broad use in behavioural clustering methods uncovered by this review. This will be discussed further in Chapter 9.

The final measures in the continuous bet domain are the three closely related behavioural measures of **maximum bet size**, **minimum bet size**, and **bet size range**, which capture in simple terms the scope of a player's wagering. Each of these three measures have only been used once in the literature gathered by the systematic review, making it difficult to determine their applicability both between existing domains, and with respect to the emerging domain of decentralised gambling. For this reason, discussion of these three measures is not taken further, although further studies in the field may provide greater context for their application in new data sets.

Discrete Bet Domain

The measure of **amount wagered per day** can be described as a measure of centrality of the sum of bets on each active betting day. Like other centrality type measures, it can therefore be computed at the aggregate level if parametric variant is used. Amount wagered per day has seen three

uses, first by Percy et al [72], then by Han et al [107], and finally by Auer and Griffiths [76]. These three studies each look at very different topics, specifically machine learning, identification of roles in illegal online gambling, and the impact of the COVID-19 pandemic on casino gamblers. Only one of these (Percy et al) use harm related variables, making it difficult to conclusively relate this measure to such variables. For this reason this measure is not discussed further.

Unlike the ‘raw’ total amount wagered described above, the **total amount wagered across duration** is simply the total amount wagered divided by the duration. This intuitively yields a number which describes the amount wagered across a player’s career were they to gamble every day of that career. This measure was first used by Fiedler in his 2011 study on poker players [65], but has not seen broader adoption in the studies returned by the systematic review performed as part of this thesis. Similarly, the measure of **variance of total amount wagered per day** has only been used once, and like amount wagered per day above this was by Han et al in their exploration of roles in illegal online gambling [107]. The limited uses of all measures in the discrete bet domain make them unsuitable as candidates for developing our understanding of players in new domains, but may be a fruitful area of further work in existing and labelled domains.

Session Bet Domain

Both behavioural measures in the session bet domain have seen only a single use in the studies returned by this review. This includes the **amount wagered per session**, and the **variance of amount wagered per session**. Interestingly, the former was first used by LaPlante et al in their 2009 study of poker players [100], but has not seen use since, despite appearing to yield notably different distributions between heavily involved and non-heavily involved groups of players. Although it should be noted that this may be due simply to the way in which heavily involved players were distinguished (by total amount wagered in the continuous domain), which likely correlates strongly to amount wagered per session. The measure of the variance of amount wagered per session has only been used once by Finkenwirth et al [101], who found that it was highly important in a random forest model of classification of self-exclusion status in online gamblers. Unfortunately it’s limited use means it may not find meaningful use in new domains, although

these recent (2021) results indicate that it could become an important measure to use in future work. Like the discrete bet domain before, the session bet domain is somewhat underdeveloped and therefore not of further use in the context of this thesis.

5.2.3 Loss Domain

Behavioural measures in the loss domain are unique in that they require additional transaction data not used by any of the measures discussed in previous sections. Specifically, they require the resulting payout from a given bet, such that the player's loss (or gain) can be computed for each bet placed. Formally, this can be represented using the finite sequence of payouts P , where $P = (p_0, p_1, p_2, \dots, p_n)$. This sequence can have a maximum length equal to the number of bets placed, but may contain zero or null values in cases where bets were unsuccessful. The sequence P may also be shorter than its counterpart B , as bets whose outcome is unsuccessful may simply not be included in the sequence. Additionally, in cases where no bets were successful, the sequence P may be of length zero. In order to provide a consistent representation of payout sequences in this thesis, the sequence P will always be padded with zeros for unsuccessful bets, resulting in the sequence always being of equal length to B .

The use of payout data unlocks a number of measures which can be used to describe a player's activity, although not all of them contribute more information to what is already known about the player from the measures described above. For example, the duration computed across bets and the duration computed across payouts would likely not yield measures which are meaningfully different. As with the core sequences in the time and bet domains, the sequence P can be divided into discrete and session domains.

Continuous Loss Domain

Like the two continuous domains above, the continuous loss domain treats all payouts as a single sequence. Two behavioural measures in the continuous loss domain are particularly prevalent in current literature; these are net loss, and percent loss. **Net Loss** is simply the sum of all payouts minus the sum of all bets, and like total amount wagered above has been used over 30 times. Here, since the measure takes bet sizes from payout sizes, a positive net

loss indicates an actual loss of value, whereas a negative net loss indicates an overall gain. As the house edge is ultimately taken from payouts, this derivation of net loss is functionally identical to what several authors refer to as the gross gaming revenue [84], which is an industry centric term referring to the amount of revenue generated by a given player as a result of their gambling activity. In the remainder of this thesis the term net loss is used, as in the context of transaction data in isolation it more accurately expresses the computation, but the two are effectively interchangeable.

Net loss has been positively associated with a number of harm related variables, including account closure due to gambling related problems [92], subsequent online gambling [108], and self exclusion [93], although many studies do not explore its potency in relation to external variables, but rather as a descriptive measure for groups of players. This includes many population level analyses [61][2][91][100], indicating that net loss as a behavioural measure is be as important for describing population level economic effects as it is a measure to associate with gambling related harms. Additionally, few measures uncovered in this review have found such broad application across a range of gambling domains. This makes net loss an ideal candidate for application in new domains.

In addition to net loss, **percent loss** is the net loss represented as a fraction of the total amount wagered. Since the measure of net loss is simply the difference between the sum of bets and payouts, it can be a maximum of the total amount wagered multiplied by the odds (if every bet is successful), and can be a minimum of the negative total amount wagered. Similarly, percent loss can reach a maximum value of 100%, and has a minimum value equal to the result of the sum of P as a percentage of B if each bet was successful in a given game. Despite requiring identical information to compute as net loss, percent loss is less commonly used, with a total of 12 uses in the studies returned by this review. Unlike net loss above, percent loss has been used almost exclusively in studies which describe populations of players, rather than exploring its relationship to harm related variables. For this reason, percent loss is a strong candidate for providing population level descriptions in a comparative way, but lacks concrete links to harm related variables so is not discussed in this context further.

The final continuous loss domain measure to be discussed is that of **sum of payouts**, which as the name implies is simply the sum of the sequence

P. First used by LaBrie and Shaffer in 2011 [94] this measure, also known as ‘Total Winnings’, has been used a total of 5 times in work by Gainsbury [103], Leino et al [90], Luquiens et al [66], and Auer and Griffiths [109]. These studies vary in their motivations, from developing predictive methods for limit-setting [109], to population level exploratory work [90]. At the population level, sum of payouts was not found to be a significantly different measure between self excluders due to addiction motives versus commercial reasons [66]. Similarly, in Auer and Griffiths’ predictive work, their variable of ‘Amount Won’ only appears important (number 5 out of 5) to one of the five predictive methods employed [109]. This lack of a relationship between sum of payouts and any harm related variable means it is not considered especially useful in the context of this thesis so is not discussed further.

The remainder of the behavioural measures in the continuous loss domain have each only been used once. These include the mean/median of the sequence *P*, known as **payout size** [72], the **payout-bet size ratio** [53], the presence of a **big win** or **relative big win** in a player’s transaction sequence [97], two clamped variables (clamped net win [101] and overall loser (binary) [110]), and finally the presence of **sawtooth occurrences** within the transaction sequence [106]. The single use of each of these measures makes their application to new domains difficult, as any findings have not been replicated and their experimental domains and configurations are broadly variable. Furthermore, computation of the final measure of sawtooth occurrences could not be replicated even if it were applicable to new domains, as the original code used and a granular enough description of that code is no longer available².

An additional computation on player transaction data which meets the definition of behavioural measure used in this thesis, is that of **payout-bet count ratio**, alternatively described as ‘the average hit frequency of a game’ by Leino et al [90]. In Leino et al’s study it is used to describe game characteristics, but it has also been used by Ukhov et al [53] - although was found not to be useful for distinguishing problem gamblers in casino or sports betting. The game centric nature of this measure and its limited use means it will not be discussed further.

²Discovered following correspondence with the original analyst on the project.

Discrete Loss Domain

Only one measure exists in the discrete loss domain; **net loss on last day**, which was used once by Kainulainen in 2021 [110]. He found that, using binary variants of net loss on last day, that this measure is a significant and positive predictor of the amount of time to the next betting day. In addition, this effect was found to be weaker in gamblers with higher experience, and that older gamblers tended to return sooner after losses than younger gamblers. These findings may be important for building a more holistic understanding how sequences of losses (or wins) affect future play, but in the context of this thesis may not be useful as they focus on player analysis at the individual/session rather than population level.

Session Loss Domain

As with the discrete loss domain above, the session loss domain has just a single measure, which has only been used once across all of the studies returned by the systematic review. The measure of **net loss per session** was used in Finkenwirth et al's 2020 study on predicting self-exclusion [101], but was not presented in their results section so is assumed to be of little importance to their aim. For this reason, this measure, and indeed this subdomain, are not discussed further.

5.2.4 Risk Domain

The fourth and final distinct behavioural domain is that of the risk (in the probabilistic sense) taken by the player. Concretely, this domain concerns any variables derived from the finite sequence O , where $O = (o_0, o_1, o_2, \dots, o_n)$. Here, each o_x holds the decimal odds for a given bet, completing the set of information available for each bet described in Section 4.3.2 above.

Surprisingly, the only behavioural measure which uses any odds data is that of **mean/median odds per bet**, used in Xuan and Shaffer's study of behaviour prior to account closure [92]. Despite its early (2009) use, none of the subsequent studies in player tracking research have applied it, which may be due to the more specific data on each of the bets that it requires. It is therefore unclear how generalisable its application can be, although Xuan and Shaffer's work finds that self-identified problem gamblers seek shorter (less risky) odds than the population in the days preceding account closure.

This potentially important finding needs further examination, although first in the context of a single domain rather than for understanding new domains as presented within this thesis. For this reason, the risk domain is highlighted as being particularly underutilised in comparison to the other domains, but is not discussed further.

5.2.5 Miscellaneous Domain

Not all of the behavioural measures uncovered by this systematic review fit neatly into one of the three domains described above. While they can still be plotted in the atlas of behavioural measures (see Figure 4.2), their allocation to the ‘Miscellaneous Domain’ may be due to their incorporation of an additional data source in their calculation, or that they are specific to a certain type of gambling activity - for example, the measure of **total rake paid** in poker [65][63].

The most widely used behavioural measure in this miscellaneous domain is that of theoretical loss, first proposed by Auer et al in 2013 [75]. This measure is simply the product of the house advantage and bet size for each bet placed, and has been subject to lengthy debate across a number of papers [75][67][84][111]. The core of this debate is whether or not total amount wagered or theoretical loss is a ‘better’ measure of gambling involvement (or gambling intensity). This is of course dependent on how this term is defined, with Auer & Griffiths using ‘the amount of money that a player is willing to risk’ - which is explained by the computation of theoretical loss [112]. This debate is covered comprehensively in Chagas & Gomes’ review [1], so is not presented in detail here, although to summarize Chagas & Gomes pose that ‘the circumstances must influence the methods and tools chosen’. In the context of this thesis, this means that in cases where house edge data is available and comparable studies have been done, theoretical loss may be applied. Unfortunately, of the 8 papers which present an application of theoretical loss, only two relate the measurement to a harm related variable (limit setting [109][75]). The remaining 6 papers each explored the effects of personalised feedback and other feedback mechanisms in gambling operators, so are not pertinent to this thesis [67][112][85][84][86][73]. This limited number of comparable uses means theoretical loss is arguably less meaningful for applications in new domains, so is not discussed further.

Several other behavioural measures discovered either do not fit neatly

into any of the above domains, or apply some external context to their computation. For example, the measure of **portion of late night bets** and **portion of bets placed on a Saturday** used by Ukhov et al [53] exist in the continuous time domain but do not strictly use all of the values in the sequence. Similarly, the **ratio of total amount wagered on weekends** and the **ratio of the number of bets placed on weekends** used by Braverman et al [69] use specific days/times as cutoffs in their computation. Each of these four measures has been used once in the studies returned by this review, so are not discussed further, although it is important to note that the concept of weekend or late-night bets does require the additional context of the timezone in which the bets exist. In the context of this thesis, and in the cryptocurrency domain in general, this is not always possible.

Discussion of each of the unique behavioural measures uncovered by this systematic review is now complete. Before moving out of the gambling literature and into the cryptocurrency domain, a brief discussion around how specific groups of measures have been used to identify behavioural profiles of players is required. This will ultimately inform how the measures above can be applied to the cryptocurrency domain, and which behavioural groups exist in comparable domains.

5.3 Behavioural Measure Clustering

A significant limitation of all of the behavioural measures work above is that any single measure cannot fully capture the nuances of a person's gambling behaviour. This is to be expected, as in this context each new behavioural measure computed simply adds a new descriptive dimension to the behavioural profile of a player. Taking an example from a different domain; when understanding different ability levels in sports, it would not be very meaningful to describe different players according to just their height, weight, or combination of the two. Instead, a richer set of measurements can be taken to describe each player in a more granular way, such as their time spent playing, training, number of games won, etc.

This example, as in clustering behavioural measures in gambling studies, raises questions around which measurements to take for each observation in the data set. This is especially important when considering classifications as potential problem gamblers, at risk gamblers, and recreational gamblers.

Intuitively one may wish to take every possible measurement available, feed all of them into a clustering algorithm, and see which ones explain different cluster memberships the best. This approach, while perhaps valid in the sense of understanding the underlying structure of the data in isolation, can only be more broadly meaningful if the resulting clusters can be related back to some existing ground truth. For example, using this technique, one may find that the number of different pairs of shoes worn by an athlete strongly explains their ability level. This finding is only useful if the measurement of ‘number of different pairs of shoes worn’ has been related to some external variable such as a player’s overall experience, as those who compete more often will naturally be more likely to have higher values for this measurement. One can therefore say that the number of different pairs of shoes worn is a proxy measure for overall experience, although as in player tracking the concept of overall experience (or gambling involvement) may have multiple concepts within it.

Focusing explicitly back on gambling studies then, much of the existing literature has aimed to classify different groups of players according to their behavioural profiles along a set number of behavioural measures - typically with the aim of capturing both financial and temporal aspects of gambling involvement. One of the earliest studies in this area was Braverman and Shaffer’s work identifying behavioural markers for high risk internet gambling using a player’s first month of betting data [3]. Braverman and Shaffer found that four distinct gambling behaviour profiles exist in their sample of 530 live action gamblers. These included a high activity & high variability group, a low activity group, a high activity & low variability group, and a moderate betting group which made up the majority of players. Crucially, they found that 73% of the high activity & high variability group had reported closing their accounts due to gambling-related problems. This is in contrast to 45%, 29%, and 32% of the other groups respectively. Whilst the limited sample size of this study with respect to the number of measures used (4) and the number of clusters identified (4) means one can only tentatively generalise its findings, it indicates that important factors about players can be inferred from their transaction data alone, and that clusters within such data can reveal meaningful and externally relevant player behaviours of interest.

Adding to the identification of player profiles, Braverman, working with LaPlante, Nelson, and Shaffer in 2013 [69] then employed a different machine

learning technique (CHAID decision tree rather than k-means) in order to identify characteristics that distinguished a subgroup of high-risk internet gamblers from within a sample of 4,056 players, again using their first month's betting data. Their results on this larger data set confirm, again with this different clustering method, that highly variable gambling is an important marker for the development of gambling related problems. They also note that participation in multiple gambling activities contributes to making this distinction, although in the context of this thesis' focus on transaction level data this remains an important but non-applicable finding.

The previous year, Gray, LaPlante, and Shaffer's study into the behavioural characteristics of internet gamblers who trigger corporate responsibility gambling interventions [95] found that measures of gambling intensity (total number of bets and bets per day) could be used to distinguish from controls. This adds to the idea that for describing a population of gamblers in general, a broad set of behavioural measures can be used, but when building profiles of subgroups within that population, measures which have been used before, and the methods in which they were analysed, should be replicated for the most directly comparable results. This topic of replication and compatibility is discussed as encountered in the behavioural clustering study in Chapter 9.

5.4 Cost of Gambling

Before concluding this chapter, a brief discussion of factoring in the cost of gambling into behavioural measures can be presented. This is particularly important in the context of cryptocurrency gambling which this thesis explores, as the cost of gambling can come from gas fees incurred by a transaction (See Section 3.3.2).

This systematic review has identified many behavioural measures which focus not on the broader cost associated with gambling as an activity, but rather of the costs incurred by each individual bet. This means that although individuals costs of gambling (including subscriptions to gambling platforms, travel expenses, etc) do broadly fall into what can be described as the costs associated with gambling, they have not historically been relevant to describing an individuals gambling in the context of behavioural profiling. Gas fees can then best be described as a fixed non-refundable fee paid by

each player for each bet placed. For this reason, gas fees will not be included in further studies in this thesis, as they are considered out of scope of the behavioural profiling techniques explored in this chapter, and indeed the field. This is not to say that analyses of this nature is not useful, but rather that it sits closer to the study of the economics of gambling, and not player behaviour tracking.

5.5 Summary

This chapter has introduced a sub field of gambling studies known as player behaviour tracking, and the many behavioural measures which are applied within it. Breaking down the behavioural measures uncovered by the systematic review in the previous chapter according to the information required to compute them has revealed a mix of widely-used and barely-used measures across a range of studies. This information-oriented taxonomy reflects the broader context of this thesis - specifically the data available in decentralised gambling applications - showing that many of the most widely used measures can be computed with bet time, size, and payout information. Using a systematic review, over 50 unique studies were discovered, employing a total of 51 unique behavioural measures. These measures each add a new descriptive dimension to a player, although not all of them can be concretely associated with an external harm related variable such as probability of self exclusion or self reported problem gambling severity. Furthermore, a subset of studies have been identified which aim to identify behavioural groups within populations, rather than simply describe them. These studies use more granular behavioural measures than their population-level counterparts, and provide a promising set of measures and methods to be applied in new domains.

In the context of understanding new domains such as cryptocurrency gambling, it appears a replication-oriented approach would yield the most meaningful results, as although several of the measures uncovered have seen widespread use, the number of different experimental configurations and gambling activities makes selecting a generic *and* meaningful single set of behavioural measures difficult. This limitation is inherent to the size and age of the field itself, rather than a reflection of the quality of the work therein, however, it is not severe enough to prohibit the exploration of new

domains, as later experiments in this thesis aim to show. The next chapter in this thesis begins applying metrics to actual transaction data, before the measures described here are applied in Chapters [7](#), [8](#), and [9](#).

Chapter 6

Decentralised Gambling Application Prevalence

“If your experiment needs a statistician, you need a better experiment.”

Ernest Rutherford

Chapters 2 and 3 described the fundamental concepts behind cryptocurrency networks, and the process for extracting their transactions to produce collections of useful behavioural data. Before the behavioural profiling techniques uncovered in Chapters 4 and 5 can be applied, an analysis of the prevalence of gambling *and* gaming applications in the cryptocurrency domain is required. This follows that a number of popular blockchain games operating atop the Ethereum network contain payments to randomised reward mechanisms which are mechanically similar to gambling applications [9]. Such mechanisms within blockchain gaming applications are considered a form of gambling in this thesis, and are therefore of interest in the context of understanding gambling in this emerging domain.

The motivation behind these prevalence analyses are two-fold. First, an understanding of the scale of use of cryptocurrency technology in both gambling and gaming provides essential context for subsequent behavioural profiling studies - a context unavailable in existing peer-reviewed literature. While several application ranking services do provide prevalence data, it is not clear whether the results presented by such services are based on actual blockchain transaction data, or are influenced by the presence of

misclassified applications. Second, by understanding the prevalence of such applications, a sample of market leaders can be extracted whose transactions contain a large sample for subsequent research. This is important as decoding transactions as described in 2 is a resource intensive and application-specific task, so knowing which applications may provide the most data is essential to maximising sample size in subsequent research.

6.1 Study 2: Gambling Application Prevalence

Before applying the behavioural measures described in Chapter 5 to decentralised gambling transaction data, a preliminary analysis can be performed with the focus of understanding the scale and trajectory of the currently available applications. Such analyses will help frame further studies in a broader context, and contribute to our understanding of the scale of use of this new technology. It is also important to apply the insights generated by subsequent chapters rather than simply re-print analyses which are freely available on generic blockchain explorer applications or ranking services. For example, the StateOfTheDApps ranking service provides a collection of interactive visualisations¹ which offer insights into the trajectories of a number of metrics across their set of application genres. This section therefore presents analyses specific to the data set gathered as part of this thesis.

6.1.1 Introduction

In the sparse decentralised gambling domain literature [87, 113], very little is known about the scale and overall trajectory of the set of available applications. This presents a problem for interpreting the results of such work, as results are difficult to contextualise. For example, if the average involvement per player is larger than in comparable centralised gambling, but only a few players actually use these applications, then it may be more effective for gambling research to remain focused on centralised data sets. However, if the population or expenditure of decentralised gamblers is growing, then given the immutable and globally accessible nature of these applications, it could raise significant regulatory and public health concerns at a scale greater than or equal to their centralised counterparts over the coming years. To this end,

¹Available at <https://www.stateofthedapps.com/stats>, accessed 29/04/2021.

metrics of interest include the number of users over time, the total amount of value transacted over time, and the total number of transactions over time. These three metrics can be broadly applied to any web application, and indeed any e-commerce sites, and are chosen as they have been used as precursors to discussions around behavioural measures[61][3], and in similar work understanding application growth across cryptocurrencies [49].

Using these three metrics the following research questions can be addressed, which extend research question 3 posed in Chapter 1 - How prevalent are decentralised gambling applications?;

1. What is the rate of spending in decentralised gambling applications?
2. What is the rate of new user adoption?
3. What is the volume of betting activity?

Additionally, market-leader driven questions can also be addressed, which may support the idea that a representative sample of players in the decentralised gambling domain can be generated via the analysis of a select set of market leaders, rather than requiring in-depth analysis of many applications;

1. Which decentralised gambling applications contain the most transactions?
2. Which applications have the most historical users?

6.1.2 Method

Data Gathering

The public and transparent nature of the Ethereum blockchain means every decentralised gambling application that has ever existed on the network is available for study. In order to convert this abstract truth into a list of actual applications, an application ranking service can be used. Such application ranking services function similarly to top music or film charts, but instead rank decentralised applications. In the context of this thesis, the rank given by these services is not meaningful, as all of the transactions for all of the applications will be gathered regardless. Instead, such a service is useful only for identifying the names of existing applications and the addresses of the smart contracts they use. One such service is StateOfTheDApps,

a multi-cryptocurrency, multi-category (gambling, fintech, games, social, etc) service which can be used to create a list of decentralised gambling applications and their associated smart contract addresses.

Using StateOfTheDApp.com's search functionality, a list of gambling applications on the Ethereum network can be retrieved, which returns a total of 195 unique names at the time of writing². Each application has its own page on the StateOfTheDApps service which lists the addresses of any associated smart contracts. This list, and indeed the existence of a given application on the site, follows an open submission process which can be performed by anyone, but is advertised towards developers and marketing teams behind these applications³. This study relies on the StateOfTheDApps tag system, with tags being assigned as part of the submission process by the submitting party.

Usage Metrics

In order to plot each of the three metrics described above across the applications in the collection, all of the normal transactions (see Section 3.2.1) should be retrieved, as these generally represent the actions of users rather than the internal actions of the applications themselves. Once isolated, the number of users over time can be visualised as the cumulative sum of the number of unique addresses found in the normal transaction set for each application (sorted chronologically). The first difference of the cumulative sum can also be used to show the rate at which new users are entering the decentralised gambling application ecosystem, although as discussed in the limitations below, this is not a perfect representation. The total amount of value transacted over time can be represented using the cumulative sum of Ether sent to each of the application's smart contracts. As with the number of users, the first difference of this data can be used to assess the rate at which the total value is increasing. Finally, the raw number of transactions can be represented in an identical way, again with first differencing allowing a broad understanding of the rate of change over time.

With collection-level metrics presented, the distribution of the top 20 platforms by each metric will be presented. This top selection of platforms should be similar or identical across all of the metrics, with any discrepancies

²25/04/2021

³See <https://www.stateofthedapps.com/dapps/submit/new>, accessed 25/04/2021.

e.g. an application in the top 20 by user count but not transaction count, potentially indicative of non-human activity or some other difference of architectural or phenomenological interest. This idea of suspicious activity due to botting will be discussed in more detail in the dedicated Section 7.3.3 below. These top collections present the most fruitful applications for subsequent research given their larger data sets than the majority of applications. As with other classes of web applications, the distribution of players across each of these metrics is expected to be severely skewed, with a handful of applications enjoying a majority market share. Reasons behind this skew are beyond the scope of this thesis, but how this skew compares to non-decentralised applications and indeed decentralised applications of other types poses an interesting area of future work.

Data Sample

The data set gathered for this macroscopic market analysis of decentralised gambling applications included 13,956,372 transactions across 147 unique applications. This number of applications is lower than the 195 above as some may have no transactions, only internal transactions, or may not have contract information entered correctly or completely in their StateOfTheDApps entry. These transactions span from 07/06/2016 through to 26/07/2021, or block number 1,660,887 to 12,900,878 on the Ethereum blockchain, providing an extensive longitudinal sample across world events like the COVID pandemic, price bubbles and crashes, and the introduction of new applications to the ecosystem.

6.1.3 Results

Population Level Metrics

The first population level metric of interest is the total value transacted over time across all applications in the sample. Figure 6.1a visualises this sum, totalling almost 5,000,000 ETH by mid 2021, equivalent to almost 21.5 billion dollars with a current ETH price of \$4,300⁴. This static current valuation is naturally not wholly meaningful in the context of gambling data, as payouts of successful bets which are re-spent all add to the total, but nevertheless it can be used to compare with centralised gambling operator metrics, or

⁴See <https://coinmarketcap.com/currencies/ethereum/>, accessed 01/11/2021.

used with house edge percentages to compute approximate gross gambling yield across the market. With this in mind, the plot shows the year 2019 was significant in terms of value growth in the ecosystem, increasing almost 600% from 0.5M to 3M ETH. A further almost 50% growth to 4.3M ETH in 2020 brings us close to current levels. An interesting artefact of the cumulative value transacted is that it does not appear to have been noticeably affected by the COVID pandemic - exploring this further is considered out of scope of this study, but presents an interesting comparative opportunity to similar work in the centralised domain [76].

The weekly change (first difference) by its nature shows the exact rate of change in the cumulative sum, highlighting a strong peak in the second quarter of 2019. These do not appear to correspond with any significant changes in other metrics, nor do they visually map to any drastic change in price (See Figure 2.5 in Chapter 3). Instead these may be artefacts of the launch of new applications, the result of botting activity, or some other external effect. Both the cumulative value and the first difference in cumulative value are particularly meaningful in the context of the price of the underlying cryptocurrency, Ether. With a sharp increase from late 2020 through to 2021, although the raw weekly change in cumulative value (ETH) appears low historically, in the context of underlying growth from £500 to almost £3,000 in the same period the usage of decentralised gambling applications in real terms appears to be increasing through 2021. This may be due to the increase in user numbers in 2021, as discussed next.

The second population level metric of interest is that of ‘player’ count, or in blockchain terms, the number of unique addresses transacted with over time. Similarly to the cumulative value above, the number of unique addresses shown in Figure 6.1e has seen growth across the time period, although this time peaking in late 2019. In the context of a total 5M ETH spent in Figure 6.1a, the additional context of the number of unique addresses shows that applications saw relatively high rates of adoption as early as late 2017 although this early adoption did not lead to a spending increase. Interestingly, this peak in 2018 does coincide with a price peak of \$1000 at the start of 2018, although given the nature of such comparisons this may simply be a coincidence.

The change in unique addresses, as with change in cumulative value, appears extremely variable with no obvious trend. Notable features include a

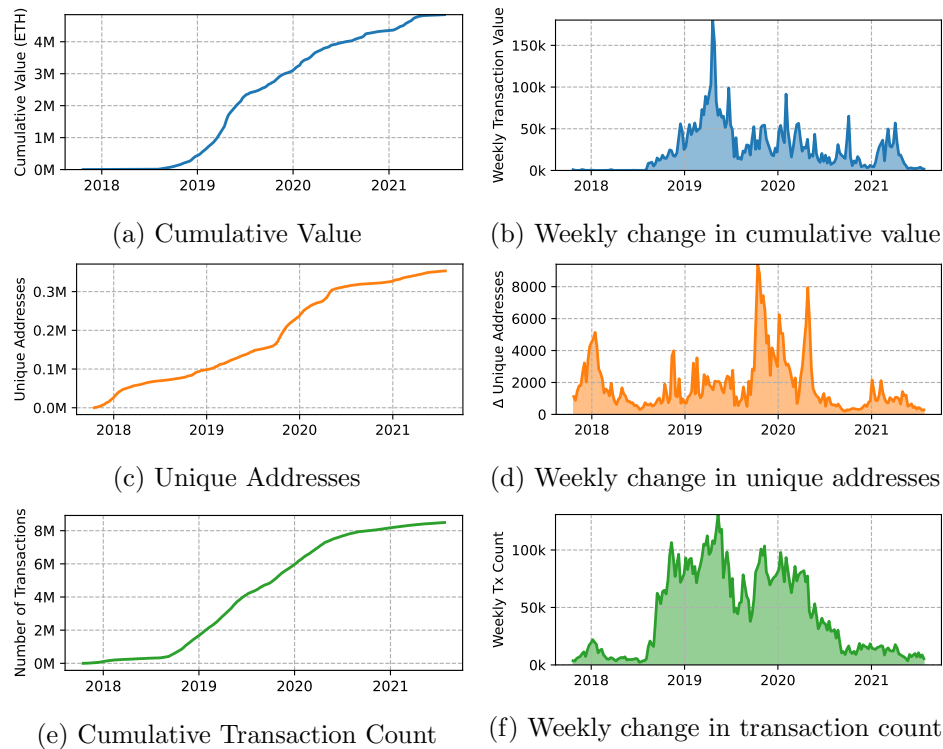


Figure 6.1: Usage metrics across all 147 decentralised gambling applications gathered as part of this thesis.

period of sustained and increasing uptake in across the 2019 winter, although this abruptly ends, resolving to a gentler increase across 2020 through 2021. This 2019 spike is an area of phenomenological interest as it does not appear as strongly in any of the other metrics. This may indicate the explosive growth of a single application or group of applications following some viral marketing strategy, or may be the historical artefact of foul play such as the introduction of non-human players to the ecosystem. This is explored in more detail in the following study in this thesis.

The final population level metric is that of cumulative transaction count. In the raw transaction data context, it cannot be assumed that each of these transactions equate to the placement of a bet, but rather a financial interaction with an application of some kind. This third descriptive dimension supports the idea of a general increase in adoption throughout 2019, and also highlights a drop in weekly transaction count in late 2020 which has not risen since. The addition of this third metric to the previous two also suggests

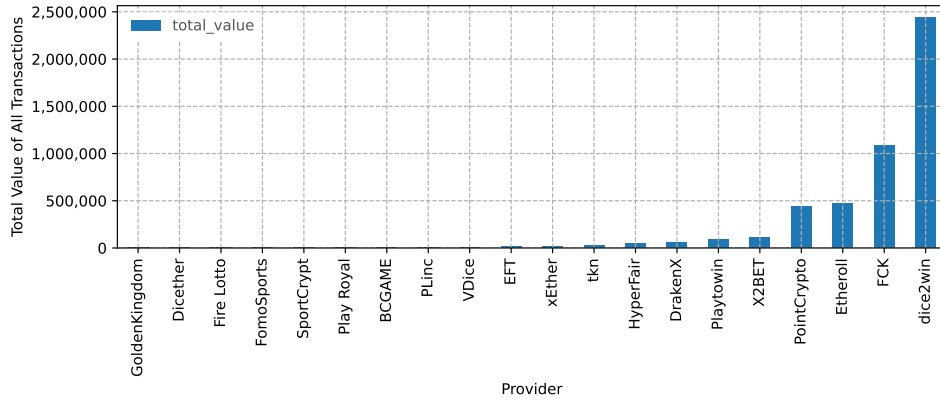
that the 2019 spike in cumulative value is not actually representative of a general increase in the number of transactions sent to these applications, as a similar peak of over 100,000 weekly transactions exists in late 2018. It is impossible to tell from these aggregate charts whether this was caused by an increase in stakes by the same players throughout this period, the creation of new applications, or any number of other causes, although it is important to remember that throughout 2021 the price of Ether has increased dramatically.

Market Leaders


With population level metrics presented, a brief review of the market leaders in the decentralised gambling domain can be presented. The metrics used are identical to those used at the population level, but are instead presented as aggregate sums for only the top 20 per metric rather than accumulations over time. Identifying market leaders in this way acts both to provide a broader context for subsequent research, and presents fruitful options for analysis as more prominent applications will likely yield more usable data once decoded.

All of the metrics describing the market leaders in Figures [A.1b](#), [6.4c](#), and [6.4a](#), are heavily skewed, with the `dice2win` application enjoying a disproportionately large market share across both the value and transaction count metrics. Both the `FCK` and `Etheroll` applications place second and third in these metrics respectively, although these top three are not leaders by unique address counts, but do appear in the top 20. Interestingly, the leaders by unique address counts (`FunFair` and `YOLOrekt`) do not appear in the top 20 by total value. This may be caused by differences in architectures and function, as the `FunFair` application offers casino games, but also offers a wallet for integrating with other decentralised applications. Similarly, although `YOLOrekt` is tagged as gambling, it provides users with a gamified options trading platform rather than casino games. These subtle differences highlight the somewhat generic nature of the `StateOfTheDApps` tagging system, although the `dice2win`, `FCK`, and `Etheroll` market leaders actually contain bet placement functionality for casino games, so present themselves as strong candidates for further study, despite ranking lower by unique address (player) count.

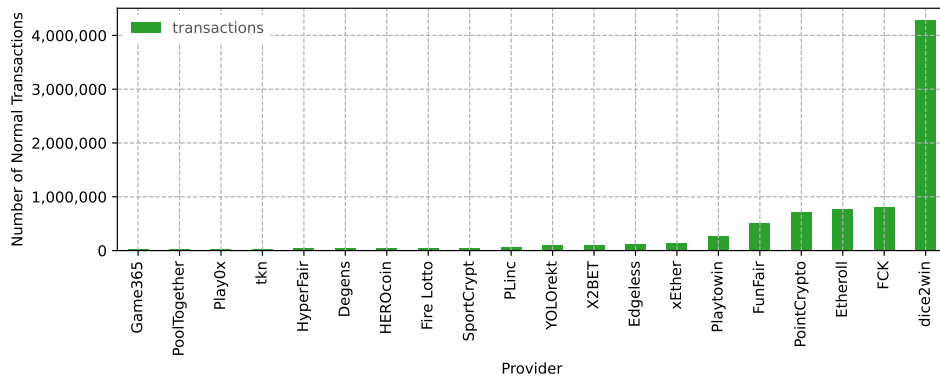
Before discussing these results and the hypotheses tested by this study, it



(a) Total Value



(b) Unique Addresses (Player) Counts



(c) Transaction Counts

Figure 6.2: Usage metrics of the top 20 decentralised gambling applications in the data sample gathered for this study.

is important to acknowledge that the results presented are in fact a second analysis of this data set, with the first discarded as all metrics were heavily skewed by two misclassified applications. Specifically, the Etherpromoswin application and the CoinGathernator applications. Etherpromoswin does not appear to be a decentralised casino at all, but rather a tokenisation contract based on ETH⁵. Similarly, the CoinGathernator application does not appear to contain decentralised casino functionality either, with the Etherscan platform reporting it as being involved in the ‘Plus Token Ponzi Scam’ - a coordinated international Ponzi scheme which saw a total of 89 members arrested and \$5.7 billion worth of cryptocurrency taken from approximately 2 million people⁶. These two applications were originally included in this analysis, but were discovered to be skewing all metrics as they appeared as market leaders despite containing no gambling functionality, so were removed. The figures generated by this first analysis are presented in Appendix A for completeness, and show how inclusion of these two application’s transactions affected the metrics.

6.1.4 Discussion

Both the population level and market leader usage metrics presented in this study offer insight into the way in which decentralised gambling applications have developed over recent years. As with the adoption of cryptocurrency technology in general, they experienced slow initial growth, and exhibit some artefacts in their growth which echo the price of the underlying cryptocurrency on which they are built. This considered, there does appear to be a disconnect in overall trends between the price chart and each of the metrics presented. This supports the idea that the scale and growth of decentralised gambling applications is not simply a function of the price of the underlying currency, but rather a distinct phenomena which grows in sometimes related but sometimes unpredictable ways. The application of more sophisticated analytical techniques to quantify the price correlation with different usage metrics in the decentralised gambling domain present an interesting area of future work.

⁵See <https://weth.io/>, accessed 28/11/2021.

⁶See <https://uk.finance.yahoo.com/news/blockchain-bites-plus-token-ponzi-164410318.html>, and <https://www.coindesk.com/policy/2020/07/30/police-arrest-27-alleged-masterminds-behind-57b-plus-token-crypto-scam/>, accessed 28/11/2021.

The results of an initial market leader analysis uncovered the Ether-promoswin and CoinGathernator applications, which are not decentralised gambling applications, but which nonetheless appeared in the top three market leaders by each metric. Similarly, PointCrypto is less of a decentralised gambling application in particular but rather a platform upon which decentralised gambling applications could exist. This finding suggests that while the methodology of taking application listings from a decentralised application ranking service does indeed uncover applications of the desired type (dice2win, Etheroll, FCK), such samples cannot be considered free from misclassifications. The large and growing number of applications invites computational approaches for smart contract classification, a problem which has already been approached by Tian et al [114] and others [115][116]. Such automated approaches could be used in the context of gambling research not only to save time removing misclassified applications in existing data sets, but to automatically search the Ethereum blockchain for applications which do not appear on particular ranking services. The key insight from this work however, is that raw ranking service categories and other tags cannot be used as training data without some level of cleaning first. In the experimental configuration used in this study, removing the two misclassified applications and re-running the analysis could completely negate their effects on the metrics, however in more computationally oriented analyses such as the machine learning approaches presented in [114], it may not always so clear from outputs that the data is being heavily skewed.

Somewhat counter-intuitively, the relatively steady growth of the three population level metrics matches examples of the rise of centralised internet gambling, with Koivula et al's study of internet gambling in Finland showing a similar near-linear trend from a market share perspective [117]. Both cryptocurrencies and the internet can be described as transformative technologies with respect to how they challenge existing ways of providing gambling services, making early linear growth a potential hallmark of future ubiquity. This projection can however only ever be anecdotal, as such technological advancements have only happened a handful of times in history, making it impossible to say at this stage whether or not the results presented here indicate future market domination of cryptocurrencies for gambling. Instead, these results simply show that the decentralised gambling domain has grown from nothing to almost 300k unique addresses having spent a cumulative 5M

ETH over the past four years on the Ethereum network alone, warranting further research in this emerging domain.

The population level hypotheses outlined in Section 6.1.1 can all be firmly rejected, as the growth across each of the metrics at the population level is clearly not exponential. Similarly, H5 can also be rejected, as although the dice2win, FCK, and Etheroll applications lead by total value and transaction count, they do not simultaneously lead by unique address count. Finally, H4 cannot be rejected, as the distribution of market leaders across each metric is indeed heavily skewed. Confirming this heavy skew supports the idea that given the high time cost of analysing application architectures, researchers will be most efficient in terms of extracting the most amount of usable data by focusing on a select few market leaders. This appears obvious, however the market leaders as presented by ranking services may not actually be decentralised casinos but rather tokenisation platforms, Ponzi schemes, or other misclassified applications. This highlights the dangers of using uncleaned aggregate data, and shows that just because all historical data is preserved on the Ethereum blockchain does not mean that it is all meaningful and usable.

Limitations

Understanding the adoption of applications which use a new technology is difficult, even when transactions to these applications are publicly available. This is in part due to the pace of innovation, with new applications appearing, and existing ones going offline with relative frequency. While applications going offline is less of a problem in the context of decentralised applications (as transactions remain available on the blockchain), un-ranked applications cannot be accounted for in the analysis presented here and are simply a known unknown. This applies not only to unknown applications atop the Ethereum network, but to applications on other networks whose transactions were naturally not a part of the data set. With this considered, the findings from this preliminary analysis provide insight into just a fraction of the decentralised gambling applications available atop cryptocurrencies in general, but likely capture a large portion of those operating atop the Ethereum network. Similarly, of the applications included in this study, at least two (Etherpromoswin and CoinGathernator) were not in fact decentralised gambling applications in the casino-game sense, but rather a tokenisation

platform and an alleged part of a Ponzi scheme respectively. This deliberate or inadvertent misclassification means that the metrics presented are at least partly skewed by the presence of similar applications, although the large number of applications gathered means this skew should not prohibit understanding of the overall market growth.

A further limitation of this relatively simplistic analysis is the potential for the presence of bots to skew the findings. As described above, this has a dedicated section in this chapter (Section 7.3.3), but it's important to note that the presence (or lack of) bots in a cryptocurrency transaction data set cannot easily be determined using macroscopic metrics. As presented in Section 7.3.3, there are techniques for identifying suspicious addresses in the context of gaming and gambling applications, so before these techniques are applied, findings may be adversely affected by the presence of non-human transactions.

A third limitation of applying these findings to understanding the scale and trajectory of decentralised gambling applications at the macroscopic level is that they do not necessarily describe the scales of harm associated with this new domain. For example, while an increasing number of users by the metric above may indicate a growing application user base, it may be the case that each of these users uses the application infrequently or only engages a few times before leaving permanently. Conversely, heavily involved users may counterbalance these lightly involved users, making an application appear fairly typical by these metrics while hiding the potential for significant harm. These metrics may therefore present a somewhat simplistic view of actual gambling behaviour. It's also important to note that the aim of this preliminary analysis was not to present an overview of the scales of harm associated with these applications, but rather to provide context for such analysis, so this limitation is overcome by the application of more sensitive and granular measures.

Finally, there is a possibility for users of the application studied above to use multiple accounts to transact with the same application. Whilst it is not clear what the benefit for doing this might be, it is nonetheless still possible and something that researchers may wish to account for in further studies. It was not considered here as to determine whether a single individual is responsible for multiple accounts, and which accounts they are, may quickly become a very involved task.

6.1.5 Conclusion

The cryptocurrency ecosystem, and the collection of decentralised gambling applications within it, have experienced huge growth in the years since their inception. This growth has been most noticeable in the past three years, and coincides with an increase in public awareness and discourse around cryptocurrencies in general [118][119]. This growth however, does not appear to have penetrated into the decentralised gambling domain, with the metrics of spending, new user accounts, and transaction volume, exhibiting approximately linear growth. This considered, several decentralised gambling applications present themselves as strong candidates for further research given their high total values and transaction counts. These are the dice2win, Etheroll, and FCK applications, although by the metric of unique address counts several other candidates are available. The heavily skewed nature of the usage metrics in this emerging domain means that researchers will benefit most in terms of extracting representative samples of players by analysing the architectures of market leading applications.

6.2 Study 3: Blockchain Game Prevalence

As discussed at the start of this chapter, blockchain games can contain randomised reward mechanisms which are mechanically similar to the simple casino games provided by decentralised gambling applications [9]. As in the previous study, before behavioural measures can be applied to blockchain games, it is first necessary to understand the broader properties of the ecosystem in which blockchain games exist. This contributes to an understanding of player behaviours in this domain at the macroscopic level, and therefore lays the foundation for more specific analyses. This can be done, like in the gambling domain, by exploring the number of users, the number of transactions, and the volume transferred to all of the applications in the data set described in Section 6.2.2 below.

6.2.1 Introduction

Much like the decentralised gambling domain discussed in Chapter 7, very little work exists in understanding player behaviours in the blockchain games domain. This may be in part due to the youth of the technology, and may also

be driven by the somewhat high technical barrier to entry that blockchain technology, and the Ethereum network in particular, incur. The development of blockchain explorers like Etherscan have helped lower this barrier, making it easier to access transactions by providing independent APIs, although broader knowledge and use of such explorers remains limited.

This study aims to establish an understanding of the rate of growth and prevalence of blockchain games. Similar work has already been done by Min et al, published in 2019 [49], although the very early nature of their work means that their analysis only extends up to late 2018. As many of the transactions analysed here fall in 2019, 2020, and 2021, a more recent understanding of the development of the domain is required. Additionally, while Min et al analyse a broad range of games, across a number of cryptocurrency networks, their analysis includes ‘Trade & Investment’ and other types of application. This study therefore focuses exclusively on self identified blockchain games (as found on the StateOfTheDApps ranking service), and those operating atop the Ethereum network given this thesis’ Ethereum-centric focus.

As in Study 6.1, using the three metrics above, the following research questions can be addressed;

1. What is the rate of spending in blockchain games?
2. What is the rate of new user adoption?
3. What is the volume of betting activity?

A similar analysis of market leaders presents the following research questions, with the ultimate goal of establishing a shortlist of applications whose architectural analysis would yield the most usable data for further research;

1. Which blockchain games contain the most transactions?
2. Which have the most historical users?

6.2.2 Method

In order to understand the development of the blockchain games ecosystem, three key metrics are applied. The first is the total value (in ETH) of transactions. This first metric may depart from previous analyses in the decentralised gambling domain, as the value components of each of the

transactions in blockchain games can be zero, with the underlying ‘value’ transfer being a token of some kind. This difference means for example that the value of transactions at this level, measured in this way, captures the expenditure from each unique address to each application, rather than the actual transfer of ‘value’ in the economic sense. For example, a transaction may have a value of 0.01 ETH which could be a flat fee set by the application itself, but within that 0.01 ETH transaction could be the transfer of a token worth 30 ETH. An example in the physical world would be using the value of containers of objects for analysis, rather than the objects themselves. In some cases (the decentralised gambling domain), the container itself represents 100% of the economic value of the objects being transferred, but here this may not always be the case. It should be noted that while the existence of NFT marketplaces makes approximate valuation of each of the objects within the transactions themselves possible, this represents a substantial task considered out of scope⁷ of this study.

The second metric is that of the total number of unique addresses transacted with. This naively captures the number of players, although as the previous study has shown, the presence of non-human players within any sample of cryptocurrency transactions should be assumed and accounted for in discussion. Despite the presence of non-human players, this metric can still be informative as it provides an upper-limit of the number of human players to be presented against the other usage metrics.

The final metric is that of the total number of transactions sent to blockchain games applications. This metric, and specifically its first difference, can be used as a proxy for the level of engagement with a given application over time. Indeed, Min et al [49] explored weekly transaction volume in both EOS and Ethereum blockchain games in their work as a measure of popularity. Additionally it is important to note that in blockchain games, each in-game action which is stored atop the Ethereum blockchain is by definition represented as a transaction. Capturing this metric may therefore not capture all actions within each of the blockchain games, but captures all of those which the application creators deemed necessary to store immutably (plus any subroutines discussed in Section 8.2.2).

As with the previous population level study in this thesis, the first (weekly)

⁷This study is framed as a precursor to further application of player tracking techniques, rather than a study of blockchain games economics in its own right.

difference of each of the metrics will also be presented. These provide a more detailed view of the weekly changes across each of the three metrics, and when plotted can be used to better visualise the rate of change, which itself may be of phenomenological or contextual interest.

Data Sample

The data sample gathered for this study includes successful transactions from 260 unique blockchain games applications. These span from 16/20/2017 through to 06/08/2021, or from block number 4,370,026 through to block number 12,970,252 in the Ethereum blockchain. This sample covers 14,564,063 individual transactions, which represent everything from in-game actions such as breeding CryptoKitties through to initialising quests in MyCryptoHeroes. Additionally, the total number of unique addresses from which transactions originate is 868,536, although as with previous studies in the decentralised gambling domain, it cannot be assumed that each of these correspond to a human player. Despite the high number of transactions in this data set, it can be loaded and analysed on a moderately equipped PC, consuming approximately 12GB of RAM for analysis using Jupyter Lab (Python), and taking 6.8GB on disk.

Gathering the set of crypto gaming transactions used in this research through the Etherscan platform took approximately 240 days on a residential internet connection with a one second delay between API requests as per their terms of service. Furthermore, as in the case of the gathering efforts in the previous chapter, one cannot simply request all transactions from/to these addresses, instead windowed requests between block numbers needed to be sent. This windowed walk along the blockchain from the creation of these applications to present is not an efficient process, and a more expensive but faster approach would be to fully synchronise an archive node and extract the data from that. Alternatively, the Etherscan platform offers a ‘CSV Export’ feature on each contract’s address page. This allows the user to manually download 5000 transactions at a time, however at the scales under inspection for this thesis this would be prohibitively slow.

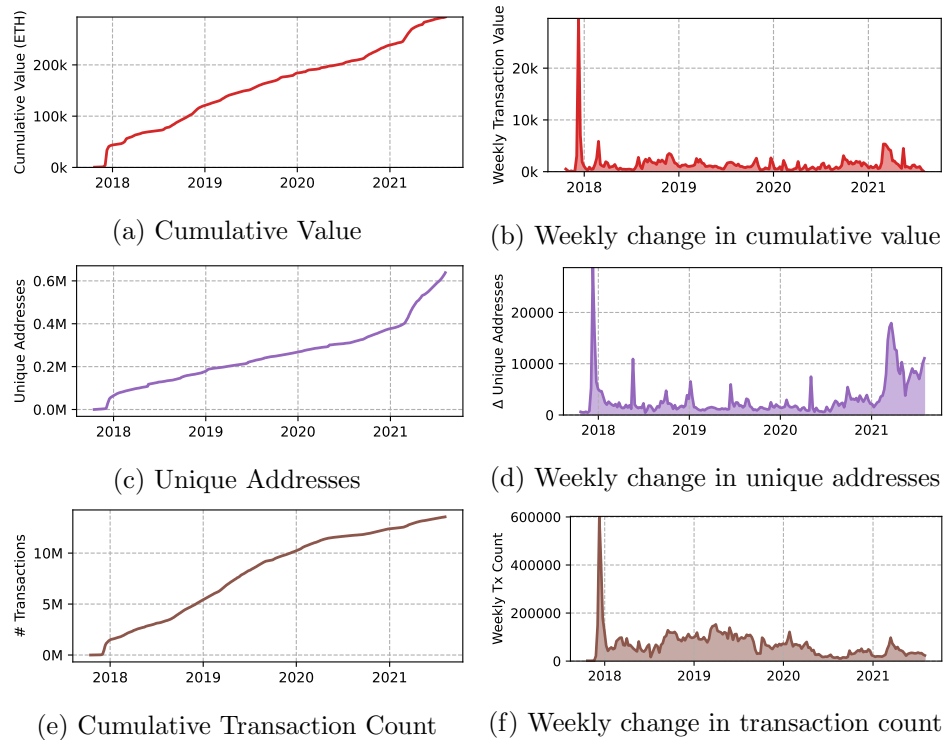


Figure 6.3: Usage metrics across all 260 blockchain games applications gathered as part of this thesis.

6.2.3 Results

Population Level Metrics

The population level metrics across each of the blockchain games gathered and their weekly first differences are presented in Figure 6.3. Unlike the same metrics in the decentralised gambling domain, the blockchain games metrics present relatively consistent growth over a period of three years starting in 2018. All three metrics begin with a sharp increase in December of 2017. This coincides with the launch of CryptoKitties, which by 10/12/2017 provides over 100,000 transactions on the Ethereum blockchain, explaining the early spikes across all metrics. Aside from these initial spikes in growth, both the cumulative value and cumulative transaction counts do not contain any noteworthy features across the data sample. This monotony is broken in the weekly change in unique addresses, which spikes sharply in early 2021 and continues to break the previous year's values until the end of the sample.

Market Leaders

An analysis of the market leaders in the blockchain games domain can help guide further research by highlighting which applications may provide the most data for subsequent analysis. Given the resource intensive process of mapping the architectures of such applications, and then using this knowledge to decode their transactions, knowing which applications to even attempt to map is critical to maximising the amount of useful data generated per amount of time spent.


Figures 6.4a, 6.4b, and 6.4c show the top 20 market leaders by cumulative value for each of the three usage metrics used in the previous section. As in the decentralised gambling domain described in the previous chapter, the blockchain games market leaders are sharply distributed across all metrics, with a top four applications enjoying the majority of the historical market share.

Unlike the market leaders in the previous chapter, the top 5 leaders in total value (CryptoKitties, Sorare, Gods Unchained, Axie Infinity, and ZED), are indeed all blockchain games, all of which focus on the collection and use of tokens in their respective smart contracts. This does not hold true across all metrics though, as the Enjin Coin application is not a crypto-game similar to those previously mentioned, but is rather an ecosystem built to support the use of NFT's created atop the Ethereum network⁸. It therefore doesn't offer any gameplay at all, but instead offers a user wallet application, several tokens (ENJ and EFI), and more. This application, like the Etherpromoswin and CoinGathernator applications in Section 7.3, has been removed from the data set and all visualisations in this study recreated without its transactions, although its presence did not influence any of the metrics as strongly as the examples in the decentralised gambling domain.


6.2.4 Discussion

At first glance, the consistent growth along all of the population level metrics presented above does not appear to provide anything of academic or phenomenological interest other than the knowledge that the growth is relatively consistent. However, in the broader context of the growing global interest in cryptocurrencies, extreme price volatility, world events, and the rate of


⁸See <https://enjin.io/>, accessed 28/11/2021.



(a) Total Value



(b) Unique Addresses (Player) Counts



(c) Transaction Counts

Figure 6.4: Usage metrics of the top 20 blockchain games in the data sample gathered for this study.

release of new blockchain games over time, such consistent growth is remarkable. Taking price volatility as an example alone; over the almost four year span of the data sample under inspection the price of ETH has varied from approximately £300 up to £3,000. Despite this, the cumulative value sent across the 260 blockchain games gathered has grown at approximately 2,000 ETH/week over the more than three year period (equivalent to £60,000 up to £6,000,000 per week over the period). This consistency (in ETH terms) shows either a growing investment in real terms per player, or the spending of existing funds by players with older funds - this is expanded on in the limitations section below.

The first research question regarding the rate of spending in the blockchain games domain can be answered as in ETH terms it is clearly linear. However, when paired with the price chart in Figure 2.5, the rate of spending in real terms appears to be increasing at a greater-than-linear rate. There are however artefacts in the price series which may frustrate efforts to neatly describe spending in real terms. For example, 2018 shows a steep decrease in price by almost 80%, 2019 shows almost no movement, and the growth in 2021 may appear somewhat exponential in the context of 2020, but taking 2021's time series alone may in fact show a linear growth in that year. The complexities of time series analyses are not expanded upon further in this thesis, but future work may employ techniques described by Chatfield to capture these artefacts in scientifically meaningful ways [120].

This study's use of cumulative unique address count has also added a non-financial dimension to understanding the rate of growth, with the number of players also growing consistently, albeit with an increase in the rate of adoption through 2021. One may tentatively map this 2021 rate change to an overall increase in public discourse around cryptocurrencies, and NFTs in particular [121][122], although it should be noted that not all NFT projects are blockchain games. As with the first research question in this study, the rate of new user adoption is increasing approximately linearly, with an apparent break in the first quarter of 2021. It should be noted however that this study only uses data from the Ethereum blockchain, hence it may be possible that the growth of the entire blockchain games ecosystem across cryptocurrencies may be greater-than-linear, but that is impossible to verify using just one blockchain's data in isolation.

The final population level question in this study, regarding the volume of

gaming activity, can again be addressed as the rate of growth by transaction count appears linear with little deviation. This is to be expected given the linear growth in cumulative unique address count, so acts as a supporting finding for the idea that the overall growth of blockchain games on the Ethereum blockchain is linear. It should however be noted that a slight decrease from late 2020 through to the end of the sample is apparent, an artefact not found in either of the previous two metrics.

The market leader based questions concerning the distribution of leaders across metrics can be answered by the graphical depiction of each metric presented in Figure 6.4. Much like the decentralised gambling domain in Section 6.1, several leaders in each metric dominate the market, or at the least have dominated the market at some stage in their existence. The second leader based hypothesis, that leaders dominate multiple domains simultaneously, cannot be so easily supported, as only the CryptoKitties application appears in the top three by each metric, with the top 10 being generally consistent albeit in a different order across metrics. This finding is not surprising given the vastly different architectures of applications in this domain - expanded upon in Section 8.2 below, as each transaction may represent a number of user actions. The key conclusion from these results in the context of this thesis is that the CryptoKitties application is a strong candidate for further analysis as it leads in total value and transaction counts, and places third by unique address count.

Limitations

The primary limitation of this study is the macroscopic nature of the metrics used. Specifically, by using only the total value, address counts, and transaction counts, more granular phenomena regarding the usage of the applications cannot be explored. For example, while the total number of unique addresses transacting with applications is broadly informative, it cannot tell us whether these addresses are engaging on a regular basis, how long they typically engage for, and other interaction behaviours. In order to explore these granular phenomena, more granular behavioural measurements are required.

In a similar vein, the analysis performed here did not warrant the use of more advanced statistical tests as used in similar work [121]. For example, it is clear from the cumulative value plot that its rate of growth is approximately

constant, hence the use of tests for seasonality, computing autocorrelation coefficients, and other time series analytical techniques are unlikely to provide any additional insight which is not obvious from the visualisation. This considered, like other studies in this thesis, computing tests used in similar work, or in similar domains, would allow direct comparison between results, however in this study this was not the objective so is instead considered an important area of future work.

A second limitation of this study is the lack of third-party data, which could be used to explore relationships between the (lack of) features in each of the population metrics versus other time series. The only third-party data gathered was the price of ETH over the observation period, which was then used as a visual comparison with each of the metrics. While this comparison enabled a deeper understanding of the underlying volatility of the price of the currency these applications accept as payment, it cannot capture important details such as when each of the users bought the ETH used in the application. For example, there is a substantial difference between a user buying ETH in 2017 when its price was roughly £300, and then spending it in 2021 when its real value has increased in value ten-fold, versus a user purchasing in 2021 and spending it in 2021. Accounting for such differences would require the use of cryptocurrency exchange data, in a similar way to the use of banking data for gambling studies [123]. This presents a challenging area of future work, but one in which important findings regarding individuals financial management strategies may be found.

A third limitation of this study is that the analysis of market leaders applied historical aggregates of each of the usage metrics, rather than an age-decayed or time sensitive metric favouring recent activity. This means that of the top applications, it is not clear whether or not they are still at the top of the market, or whether their position in these charts is due to historical positions which have since decayed. In the broader context in which this study exists however, this is not prohibitive to building on these findings, as the question of which applications may provide most data for further work is unaffected. Yet, it does mean that despite still being ‘market leaders’ in this analysis, the intensity of applications’ historical usage should be accounted for in all further work for findings to be relevant to the rapidly changing landscape of blockchain games.

Finally, this study has taken the StateOfTheDApps ‘Games’ self-identification

as a starting point for analysis. While this study confirmed that many of the top applications in the sample are indeed blockchain games (unlike the non-gambling applications in the previous study), there may still be applications within the data set which may not be considered blockchain games in the context of this study. This potential for being skewed by non blockchain games does not appear to be as significant a problem as in the previous chapter, but nonetheless persists in this domain. This can be circumvented by an initial auditing/verification step in the analysis whereby each application would be analysed before inclusion, although this would be an extremely resource intensive task as outlined in Chapter 3.

Future Work

The exact reasons for the consistent growth of the blockchain games ecosystem cannot be inferred from the transaction level analysis presented in this study alone, however, several areas of future work present themselves as a result. The most obvious in the context of this thesis is the extent to which behavioural similarities exist between spending in blockchain games and spending in decentralised gambling applications. This will involve computing more granular measures used in player tracking research such that population descriptions can be compared. An economics oriented branch of further work may apply deeper analysis to the properties of the time series of each of the metrics presented. Such analysis may explore time series modelling of each of the metrics in order to quantify trends and other features, cross-correlation with other third party data to identify relationships in growth such as social media sentiment, and so on. Finally, a comparison with different types of centralised games, gaming platforms, and other similar applications could yield useful results for projecting the growth of blockchain games over the next few years. Such speculation was considered out of scope of this study, but has important implications for regulation around this emerging domain.

6.2.5 Conclusion

This study has provided empirical support for the idea that blockchain games are steadily growing in usage, and have been doing so over the last four years. Furthermore, of the market leaders in this domain, the CryptoKitties application remains a market leader both by total value transacted and

transaction count, and appears third by unique address (player) count. This suggests that this application may yield the most data for representative analysis of player behaviour in this domain, and is of particular interest for future work.

6.3 Summary

This chapter has explored the prevalence of both decentralised gambling applications and blockchain games atop the Ethereum network. Several market leading applications in each domain present themselves as data-rich options for further analysis, although a ‘long tail’ of lesser used applications are also available. An important finding in the decentralised gambling domain is that the application ranking service provided several applications which were not correctly classified. This has implications for using such services to broadly assess the scale of spending in this emerging domain, although such applications were easily removed by manually comparing the capabilities they offered against other leading applications. The findings from both of the studies presented in this chapter lay the foundation for further research, specifically analysis of the behaviours of users of the leading applications uncovered, using the player behaviour tracking techniques described in Chapters 4 and 5.

Chapter 7

Inside Decentralised Gambling Applications

“Can’t repeat the past? Why of course you can!”

F. Scott Fitzgerald
The Great Gatsby

This chapter includes concepts and rewrites of work published in PLOS One titled ‘*Inside the decentralised casino: A longitudinal study of actual cryptocurrency gambling transactions*’.

A small portion of players who engage in gambling experience some form of gambling related harm [8]. Before applying granular behavioural profiling techniques which map to external harm related variables as discussed in Chapter 5, an understanding of player behaviours at the population level is required. This chapter therefore aims to apply all of the knowledge from Chapters 2, 3, and 5, to the domain of decentralised gambling applications. After the introductory Section 7.1, Section 7.2 presents a summary of the architectures of several leading applications identified in Study 2, describing how methods discussed in Chapter 3 could be used to generate data sets from this knowledge. This chapter concludes with the third major study of this thesis in Section 7.3, which applies a collection of behavioural measures discussed in Chapter 5 to a data set generated by the leading applications, presenting the first study of population level gambling behaviours in this new domain.

7.1 Introduction

Decentralised gambling applications are a new form of online gambling which use cryptocurrency technology to process payments and calculate game outcomes [5]. These applications vary in terms of the games they provide, and the cryptocurrencies they use. This work focuses on simple casino type games of chance, like dice rolls and coin flips, available through several applications operating atop the Ethereum cryptocurrency network. The Ethereum network is the oldest and most popular by market capitalisation of cryptocurrency networks which explicitly support smart contracts ¹, making it a strong candidate for gathering representative samples of players in this new domain. These contracts, as described in the Chapter 3, are the core technology enabling decentralised gambling applications [27].

7.2 Application Architectures

Before analysing the player behaviours in the decentralised gambling applications under investigation, it is first useful to provide broader context around how player interactions actually manifest in transaction sequences by describing their general architectures. In other words, when a player takes actions within a gambling application, which transactions are generated and are therefore which are of interest for research. This will be useful later on in the interpretation of the results, and will use several case studies to help broadly describe the entire class of decentralised gambling application online today. Importantly, this understanding of decentralised gambling application's internal architectures is the difference between being able to use their transactions for meaningful research and not, so is a central contribution of this thesis and foundational first step in the cryptocurrency gambling research domain.

Given the relative youth of the underlying technology, the core architectures of decentralised gambling applications vary dramatically, and are subject to constant innovation. However, like any other business, in order to provide games of chance to their users, a minimum set of functionality must always be implemented. These include, for example, at least one variant of bet placement function, which accepts payment from a user along with a

¹See <https://www.coinbase.com/>, accessed 12/11/2019.

collection of game parameters and executes the required random calculations before triggering a payout. Such functions may be obfuscated by different architectures, for example one may implement the ‘bet placement’ event and the setting of related parameters across a number of different function calls, or use some sort of token or other layers to complicate the core process.

Three case studies are presented below of applications whose architectures are ‘flat’ in the sense that they typically use a single contract and are therefore *relatively* easy to decode by applying the knowledge presented in Chapter 3. As with analysing any system, decentralised or not, those with more sophisticated architectures and multiple contracts will undoubtedly take more time to fully understand. This is not to say that a comprehensive architectural understanding of every platform is not possible, but rather that the amount of human effort required to do so is beyond the scope of a single thesis. For this reason, three applications have been selected based on their market position (as described in Study II above), and on their architectural complexity.

A final note before presenting the case studies is that this type of source code analysis is only possible because the original authors have submitted copies of the Solidity source code for independent verification by a third party (Etherscan). Without access to these verified submissions, a prospective researcher would need to reverse engineer the source code from the encoded blob stored atop the Ethereum blockchain. Whilst reverse engineering in general, and reverse engineering of Solidity code in particular, is entirely possible, it is a hugely time intensive process which naturally grows with the complexity of a given application. For this reason it is clear that the types of analysis performed in this thesis, and indeed developing an understanding of the architectures of these applications at all, would not be possible in any reasonable time without the efforts of the Etherscan team and the transparency of the application authors. The full source code for each of the applications below is available through their respective pages on the Etherscan blockchain explorer platform described in Section 3.1.2.

7.2.1 Case Study: dice2.win

The dice2win application presents a collection of simple casino games using a web interface (available at <https://dice2.win>), supported by a single smart contract. This application offers four distinct games, namely a coin

flip, a single dice roll, a double dice roll, and a 1–100 roll - the user interfaces for which are presented in Figure 7.2. As with other Ethereum based decentralised applications, the dice2win application relies on the browser extension MetaMask to provide the interface between a user’s cryptocurrency wallet and the application itself. In practical terms this means that in order to use this application, prospective gamblers must first acquire some Ether, transfer that Ether to their MetaMask wallet (as discussed in Section 2.3.1), and then navigate to the url above in their web browser. With these three steps complete, players can select one of the four casino games to play, set their bet sizes & chosen outcomes, and place their bets.

Contract ABI

The dice2win application’s smart contract ABI (Application Blockchain Interface) is the simplest of those audited as part of this thesis, containing a total of 20 methods. Much like different business departments in centralised casinos, gambling DApps typically include a number of different methods alongside their core game logic which perform various management, treasury, and cryptographic verification tasks. For example, the methods ‘owner’, ‘acceptNextOwner’, and ‘approveNextOwner’, each, as the names imply, relate to accessing and managing the ownership of the contract. The concept of ‘ownership’ here differs from the initial creator of the contract in the Ethereum blockchain in that it is a construct within the contract itself which is used to manage withdrawals and other privileges within the application. In practice the internal ‘owner’ of the contract, and the its original creator, can relate to the same Ethereum address, although this is naturally not a requirement. Additionally, the ‘maxProfit’, ‘setMaxProfit’, and ‘setCroupier’ functions each relate to the treasury functions of the application, allowing the owner to access or modify the contracts behaviour as programmed.

A sample list of method names is available in Table 7.1, and can be verified as an exact match to the encoded version stored on the Ethereum blockchain by comparing an identical encoding of the candidate Solidity contract to that stored in the actual blockchain. This comparison with a submitted Solidity contract has been completed as part of the Etherscan platform’s verification process for the dice2win application. This means that we can be sure that the ABI extracted, and therefore the submitted source

Table 7.1: A sample of dice2win smart contract ABI methods and their functionality (verified from address: 0xD1CEeeeeee83F8bCF3BEDad437202b6154E9F5405).

Method Name	Description
<code>acceptNextOwner</code>	Sets a previously approved ‘next owner’ (new administrator) to the contract, the owner of the contract has access to a number of executive functions
<code>placeBet</code>	Enables the placement of a bet by anyone with an Ethereum address and sufficient funds, this method is a parameterised chance function so can be used to provide multiple casino games
<code>withdrawFunds</code>	Allows the owner of the contract to withdraw funds from the contract

```

241     uint rollUnder;
242     uint mask;
243
244     if (modulo <= MAX_MASK_MODULO) {
245         // Small modulo games specify bet outcomes via bit mask.
246         // rollUnder is a number of 1 bits in this mask (population count).
247         // This magic looking formula is an efficient way to compute population
248         // count on EVM for numbers below 2**40. For detailed proof consult
249         // the dice2.win whitepaper.
250         rollUnder = ((betMask * POPCNT_MULT) & POPCNT_MASK) % POPCNT_MODULO;
251         mask = betMask;
252     } else {
253         // Larger modulus specify the right edge of half-open interval of
254         // winning bet outcomes.
255         require (betMask > 0 && betMask <= modulo, "High modulo range, betMask larger than modulo.");
256         rollUnder = betMask;
257     }

```

Figure 7.1: The dice2win application’s bitwise roll under operation, taken from source code verified by Etherscan at address 0xD1CEeeeeee83F8bCF3BEDad437202b6154E9F5405.

code, is the same as that which is in operation on the Ethereum network.

The most useful method within this ABI to this thesis’ aims is that of ‘placeBet’, which as expected handles all of the bet placement functionality for the application. This method accepts six parameters, namely the betMask, modulo, commitLastBlock, commit, r, and s. The betMask and modulo parameters contain the desired outcome of the user, and the possible number of outcomes for the chosen game respectively. These are the two most important parameters to extract as they contain actual user choices. For example, a user may be playing a single dice roll game which would require a modulo value of 6. The betMask parameter is somewhat more opaque, as it functions differently depending on the size of the modulo (See Figure 7.1). For large modulus, the betMask is equal to the ‘rollUnder’ value - the number which the internal random number generation must fall below in

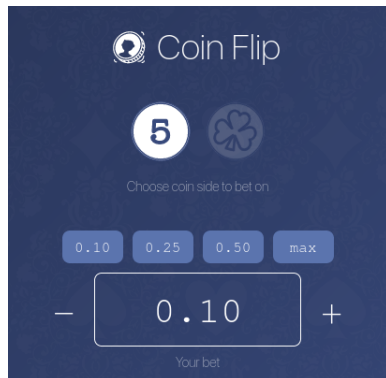
order for a bet to be won by the player (thereby triggering a payout). For small modulus it is used in the same way but via an intermediate bitwise operation which efficiently converts it into a rollUnder value. Emulating this bitwise computation is the key to unlocking the selections made by players in the dice2win application's betting functionality, and is the most important line (250) regarding this thesis in the smart contract. The remaining four parameters concern the commit value itself (breadcrumb described in 3.5.2) and the commit time-to-live, and two components (r and s) of an elliptic curve digital signature algorithm (ECDSA) [124] which are used to verify the integrity of the commit values. These final four parameters act to secure the bet placement against tampering from all parties, including malicious miners on the network. When a player places a bet through the user interface, their MetaMask wallet is used to create a transaction on the blockchain which calls the placeBet method, therefore providing each of the parameters for analysis.

7.2.2 Case Study: Etheroll.com

The Etheroll.com application is the conceptually simplest of the gambling applications whose data has been gathered as part of this research as the only game it offers is a 1–100 roll. As advertised on the application's website, *Etheroll is an Ethereum smart contract for placing bets on our provably-fair dice game using Ether with no deposits or sign-ups* (<https://etheroll.com/#/about>, accessed 13/09/2021). This description is in fact a more specific (dice oriented) summary of the application's use case, as a user chooses their desired probability of winning out of 100, bets a certain amount, and then awaits a payout from the smart contract after the underlying chance function executes. As described in Section 3.4, such a simple chance based mechanism can be directly mapped to the outcomes of large number of different casino games - including dice rolls - which makes this application in particular an excellent resource for gambling researchers.

Contract ABI

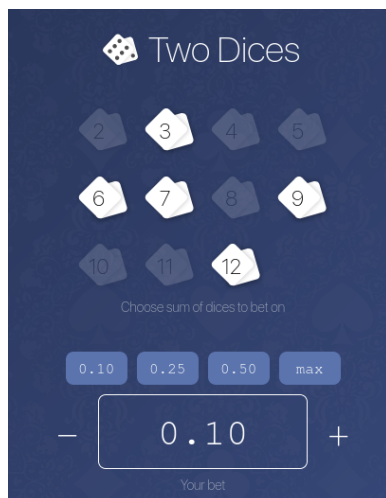
Like the dice2win application above, the Etheroll application's source code and ABI have been verified by the Etherscan source code submission process, so can be analysed with the security that the code available is true to



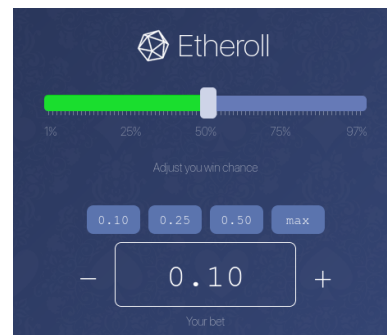
(a) Coin Flip game user interface.



(b) Single Dice Roll game user interface.



(c) Two Dice Roll game user interface.



(d) 1–100 roll game user interface.

Figure 7.2: dice2win application user interfaces available at <https://dice2win>, accessed 11/08/2021.

Table 7.2: A sample of Etheroll smart contract ABI methods (verified from address: `0xf478c8Bc5448236d52067c96F8f4C8376E62Fa8f`). Method in bold represents the core player action of interest used in further studies.

Method Name	Description
<code>ownerPauseGame</code>	Allows the owner to disable new bets, effectively closing the application to all players
<code>playerRollDice</code>	Contains all functionality required for players to play the dice roll game, including parameters chosen by the player at the time of the bet placement
<code>ownerTransferEther</code>	Allows the owner to send Ether to another address from the contract itself, much like a withdrawal but to a given address
<code>ownerChangeOwner</code>	Sets a new administrator to the contract, the owner of the contract has access to a number of executive functions including the transfer and pause methods above

that held on the blockchain. The Etheroll application offers only one game (see Figure 7.3), so should in theory be architecturally simpler than the `dice2win` application, however, this is not the case. The Etheroll ABI contains a total of 38 methods, spread broadly across a number of management, treasury, and verification functions. Unlike the `dice2win` application discussed above, Etheroll allows the internal ‘owner’ to manage the application’s parameters with finer detail, for example, implementing the ability for the owner to `ownerPausePayouts` and `ownerSetMaxProfitAsPercentOfHouse`. These finer grain controls are not essential to the function of the application, but rather allow the owner more easily modify the behaviour of the application without having to add an entirely new contract to the blockchain with the desired changes. This architecture makes the Etheroll application more adaptable to market changes such as price fluctuations, as the owner can `ownerSetMinBet` and `ownerSetCallbackGasPrice`, which set the minimum bet size, and set the gas price to be used in Oraclize methods (used as an external source of randomness within the application) accordingly.

Unlike the `dice2win` application above, the Etheroll application’s bet placement function is called ‘`playerRollDice`’, and is not a generalised chance-payout function, but contains only the functionality equal to the ‘`rollUnder`’ process within the `dice2win` application. This function accepts only one parameter (`rollUnder`), which is the value under which the random process must roll in order for a bet to be considered successful and a payout triggered. This makes analysis of the Etheroll application’s transactions more straightforward than those of the `dice2win` application above, as the



Figure 7.3: Etheroll application user interface, available at <https://etheroll.com> accessed 13/09/2021.

raw transaction data can simply be decoded and used directly without any application emulation. The disadvantage of this approach from a research perspective is that the lack of a nonce parameter makes matching payouts to bets somewhat more complex. This complexity manifests in the need to understand a ‘`__callback`’ function which exists as part of a call to an Oraclize (now renamed to Provable) function within the ‘`playerRollDice`’ call, which is a way for smart contracts to use random number data from internet sources. Understanding this function is not necessary when applying a naive chronological stitch to the bet and payout transactions discussed in Section 3.5.1.

7.2.3 Case Study: FCK.com

The final gambling application case study performed as part of this thesis is that of FCK (<https://fck.com>), which as of writing is no longer in operation. This considered, the nature of blockchain technology means that its transactions are still available, making this application of interest to researchers given its apparent large number of players, amount transacted, and relatively simple architecture.

While in operation, FCK offered an identical roster of games as dice2win, plus a roulette game. Each of the games accepted stakes in the native Ether cryptocurrency, and as with the previous case studies can be interacted with using the MetaMask wallet installed in the player's web browser. Much like the dice2win application, the FCK application smart contract employs a single chance-based method (`placeBet`) which can be called with different parameters to run and resolve bets from each of the different games. It also contains a number of management methods similar to the dice2win application, which are similar in function to those already discussed in the dice2win analysis above.

Contract ABI

The Etherscan blockchain explorer platform offers a collection of addresses linked to the FCK application. The address of most relevance in the context of understanding the key player transactions within the application is `0x9999999C60566e0a78DF17F71886333E1dACE0BAE`, to which a total of 700,307 transactions have been sent. This address contains verified source code similar to those of the two previous case studies, containing 30 methods in total, again generally divided by business function. Table 7.3 shows the full ABI of this contract.

The FCK application's `placeBet` method is functionally identical to that of the dice2win application discussed in Section 7.2.1, and includes all of the same parameters. This makes decoding its transactions an almost identical process to the dice2win application, although the bitwise operations have been slightly modified (See Figure 7.4). This exact operation differs from the dice2win application in that modulus of different sizes are handled differently through the call of a `'getRollUnder'` method. It is not clear why the contract is structured in this way over the more simple dice2win

Table 7.3: A sample of FCK smart contract ABI methods (verified from address: 0x999999C60566e0a78DF17F71886333E1dACE0BAE).

Method Name	Description
setMaxProfit	Change the maximum bet reward size, which limits the maximum possible size of bets placed to the contract
setCroupier	Change the croupier (second contract which settles bets placed by this contract)
placeBet	Allows the placement of bets to this smart contract, parameterised to provide a number of different games as an abstract chance mechanism
withdrawFunds	Administrative function which allows transfer of smart contract funds to any address

```

262     uint rollUnder;
263     uint mask;
264
265     if (modulo <= MASK_MODULO_40) {
266         // Small modulo games specify bet outcomes via bit mask.
267         // rollUnder is a number of 1 bits in this mask (population count).
268         // This magic looking formula is an efficient way to compute population
269         // count on EVM for numbers below 2**40.
270         rollUnder = ((betMask * POPCNT_MULT) & POPCNT_MASK) % POPCNT_MODULO;
271         mask = betMask;
272     } else if (modulo <= MASK_MODULO_40 * 2) {
273         rollUnder = getRollUnder(betMask, 2);
274         mask = betMask;
275     } else if (modulo == 100) {
276         require(betMask > 0 && betMask <= modulo, "High modulo range, betMask larger than modulo.");
277         rollUnder = betMask;
278     } else if (modulo <= MASK_MODULO_40 * 3) {
279         rollUnder = getRollUnder(betMask, 3);
280         mask = betMask;
281     } else if (modulo <= MASK_MODULO_40 * 4) {
282         rollUnder = getRollUnder(betMask, 4);
283         mask = betMask;
284     } else if (modulo <= MASK_MODULO_40 * 5) {
285         rollUnder = getRollUnder(betMask, 5);
286         mask = betMask;
287     } else if (modulo <= MAX_MASK_MODULO) {
288         rollUnder = getRollUnder(betMask, 6);
289         mask = betMask;
290     } else {
291         // Larger modulos specify the right edge of half-open interval of
292         // winning bet outcomes.
293         require(betMask > 0 && betMask <= modulo, "High modulo range, betMask larger than modulo.");
294         rollUnder = betMask;
295     }

```

Figure 7.4: Source code of the bitwise roll under operation found in the FCK smart contract. Source code verified by Etherscan from address 0x999999C60566e0a78DF17F71886333E1dACE0BAE.

contract, but nevertheless emulation of this functionality is required in order to extract exactly which game a player is playing, and their choice, from the raw encoded data stored on the Ethereum blockchain.

7.3 Study 4: Behavioural Distributions in Gambling Applications

The fourth major study performed in connection with this thesis concerned the application of a simple set of behavioural measures to decentralised gambling transaction data in order to establish a baseline of player behaviours in comparison to existing online casinos. As discovered in Study 1 (Chapter 5), no one set of behavioural measures can be considered perfect for describing a population of gamblers, instead a replication oriented approach is required. Furthermore, Study 2 uncovered a list of popular candidate applications whose architectures have been audited as part of the case studies in Section 7.2, allowing meaningful behavioural data to be generated from their transactions. To this end, the results of behavioural measure computation across these three applications is presented. This section contains findings presented as part of the 2020 PLOS One article by Scholten, Zendle, and Walker [87], and describes the first ever peer reviewed analysis of decentralised gambling applications.

7.3.1 Introduction

This study describes the behaviour of a large cohort of decentralised gambling application users over a 583 day period, spanning from the creation of each application's smart contracts up until the 9th March 2020 (see Figure 7.5 and 7.6). By using cryptocurrency transaction data gathered directly from the Ethereum cryptocurrency blockchain one can calculate behavioural measures using individual bet level data as opposed to aggregates of any kind, e.g. daily/weekly. Behavioural measures, including descriptions of the typical (median) player of each of the games available through each of the applications is described. Four distinct analyses were performed, following identification of likely non-human players:

- (i) a statistical comparison between human and bot players' behavioural measures




Figure 7.5: Provider-game combinations, including unique address and bet counts taken forward to the final player transaction set.

- (ii) an epidemiological description of the gambling behaviour of (human) players of decentralised gambling applications
- (iii) a statistical assessment of the relationships between existing behavioural measures of players in this new domain
- (iv) an epidemiological description of the gambling behaviours of empirically-determined heavily involved players as found in LaBrie *et al.*'s original work [2], and Nelson *et al.*'s recent (2021) study [125]

7.3.2 Method

Data Selection

Data gathered for this study includes transactions to and from three decentralised gambling applications operating atop the Ethereum cryptocurrency network. These applications were selected based both on their rank on a widely used application ranking service StateOfTheDApps, available at <https://stateofthedapps.com>, and on the subjective technical simplicity of their smart contracts. This simplicity is dependent on the author's understanding of the Solidity programming language, as encoded transactions to these contracts require decoding in order to extract the sizes of bets and player outcome selections. A deeper understanding of the language these contracts are written in, and an increase in available human resources, would increase the number of applications that could be analysed. However, given the youth of this technology, the goal here is to first understand a small sample.

Summary statistics of the data collected from these applications is presented in Table 7.5.

Table 7.4: Smart contract addresses for each decentralised gambling application used in this study.

Provider	Address
dice2.win	0xD1CEeeeeee83F8bCF3BEDad437202b6154E9F5405
Etheroll.com	0xA52e014B3f5Cc48287c2D483A3E026C32cc76E6d
FCK.com	0x999999C60566e0a78DF17F71886333E1dACE0BAE

Table 7.5: Meta data for each application gathered as part of this study. Bet and Payout values are given in ETH, and starting and ending blocks and dates represent the time window from which transactions were gathered.

	Etheroll	FCK	dice2win
Unique Users	3,086	14,466	7,868
Games	1	4	4
Bet Value	420,942.442	465,195.853	1,267,239.951
Payout Value	419,067.602	462,136.712	1,245,815.279
Start Block	6084746	6859200	6287216
End Block	9638617	8071084	9639151
Start Date	2018-08-04	2018-12-10	2018-09-07
End Date	2020-03-09	2019-07-02	2020-03-09

Data Cleaning

Transactions to and from the contracts associated with each of the applications, gathered from the start of their operations until 9th March 2020 (see Figure 7.6), yield a total of 2,232,741 bets originating from 24,234 unique addresses. Of these addresses, 14,466 transacted with the FCK application, a further 7,868 with the dice2win application, and a final 3,086 with the Etheroll application (see Figure 7.5). Figure 7.7 plots the cumulative value of the bets placed both in each application alone, and combined across the duration of this study.

The transaction data for each of these applications was gathered using the Etherscan API, which offers an interface through which transactions on the Ethereum blockchain can be directly inspected. The Etherscan API can be found at <https://etherscan.io>.

As the raw dataset is publicly available via the Ethereum blockchain, the data repository associated with this work contains the matched bets used to calculate the measures below in an accessible format (CSV). This data includes the hashes (unique identifiers) of both the bet placement and payout transactions such that the sums of the costs to and from each unique address




Figure 7.6: Transaction data gathering timelines for each of the three decentralised gambling applications studied. FCK ceased operation in July 2019.

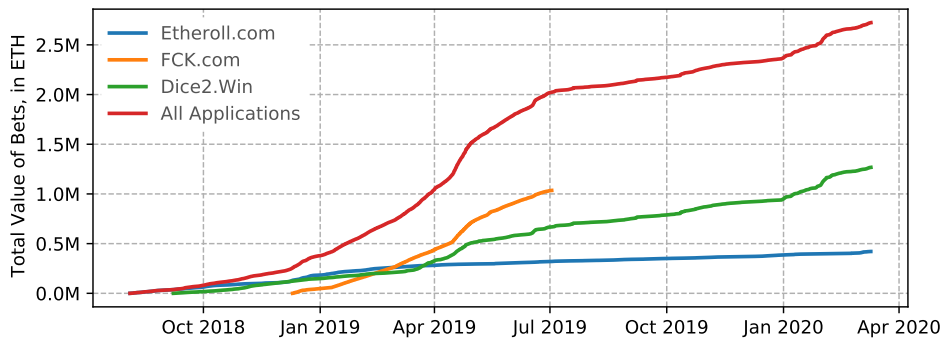


Figure 7.7: Cumulative value of bets placed through each application individually, and all applications combined over the period studied. The data and code used to create this figure is available at <https://osf.io/8bfyj/>.

can be verified. The transaction data used for this study are available in full at <https://osf.io/8bfyj/>.

Measures

The pseudo-anonymous nature of the cryptocurrency transactions from which the present data set was drawn mean that no demographic information such as age, gender, or income, is available for any of the unique cryptocurrency addresses in the set. As such, no demographic analysis was possible - this aligns with existing literature where demographic data was not found to be of particular interest in comparison to behavioural measures based on the transaction data alone [61, 2].

The variables computed as part of this study are based on those calculated by LaBrie *et al.*'s seminal investigation into internet casino games [2], given the finding that a replication oriented approach should be taken in Chapter 5. These include the duration of betting, which is calculated as the time elapsed (in days) between the placement of the first bet and the placement of the last. This is rounded up to the nearest day (following naive cast to UTC) in cases where bets were made across a midnight boundary, for example, the placement of bets both at 22:00 on a given day and again at 09:00 on the following day, are counted as having a duration of two days even though they are within 24 hours of one another. Using this I could compute the frequency of betting activity by taking the total number of days in which one or more bets was placed and dividing it by the duration of betting. This yields a percentage, with value of 100% equating to betting every day for the known duration of the use of the decentralised gambling application.

As in the original work, I calculated the average bets per day by dividing the total number of bets made by each player, by the total number of days on which a bet was placed (as used when computing the frequency above). The total amount wagered (in ETH) for each player is also retrieved, along with the total losses they incurred (also in ETH), from which their net loss is calculated. Finally, the percentage loss for each player is determined by dividing the net loss by the total amount wagered, and multiplying by 100. As in LaBrie *et al.*'s original work, the large sample size ($n=23,365$) of the players of the three decentralised gambling applications gathered in this work mean that the practical significance of any statistical differences between any of the measures calculated may be limited.

In order to promote reproducibility in this work, and to encourage further studies in this domain, the code used to calculate these measures across each of the unique addresses is available as part of the gamba library (www.gamba.dev). This library also contains methods capable of exactly replicating LaBrie *et al.*'s original work [2], plus each of the computations required to replicate all tables in the present study. The publication of the complete data set and fully documented analytical code is a core contribution of this study.

7.3.3 Results

Non-human players

Before presenting descriptive statistics for cryptocurrency gamblers, one must first ensure that the transactions used originate from human players. Given the lack of established methods in making this distinction, a naive approach, inspired by LaPlante *et al.*'s use of the Kolmogorov-Smirnov test [100], is to quantify the differences between the distributions of each of the behavioural measures for players across each of the games. I reason that if the majority of unique addresses' transactions originate from human players, collections of addresses transactions' which deviate significantly from this norm may be non-human in origin. This reasoning finds support in the fictitious scenario where an auto-betting algorithm with few parameters is used by many accounts, as this would create groups of behaviourally similar transaction sequences which would stand out. Figure 7.8 illustrates this theory, with a smaller peak indicating human players in a population with non-human players, and a second peak indicating non-human player behaviours.

To this end, I first split the collection of all gathered transactions by application, and then again by game. This resulted in 9 distinct transaction sets, each for a single application-game combination - for example; coin-flip players on the dice2.win application, two-dice players on the fck.com application, etc. The dice2.win and fck.com applications each offer 4 games, plus etheroll.com's 1-100 roll, yields 9 different games in total. From here, a two sample Kolmogorov-Smirnov test (K-S II) - which quantifies the likelihood that two samples have been drawn from the same distribution - was computed for each pair of measures, across each of the applications. This

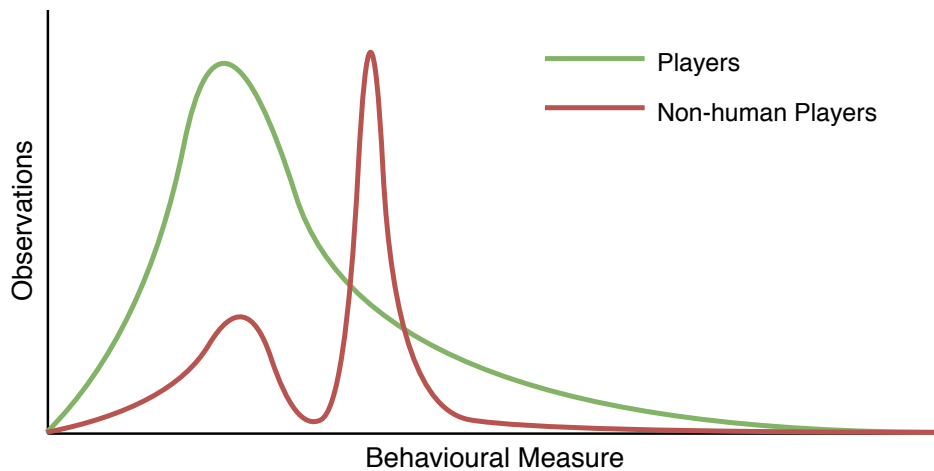


Figure 7.8: Theoretical differences in distributions of behavioural measures between human and non-human players. The second spike is created when multiple addresses transact in the same way, e.g. using by using computer code.

resulted in a $9 \times 9 \times 8$ matrix of coefficients, with axes; application-game combinations (9), application-game combinations again (9), and behavioural measures (8). Algorithm 1 shows the design of this pairwise behavioural measure comparison, a Python implementation of which is available at www.gamba.dev. It should be noted that performing this many K-S II tests without correction limits their individual descriptive power. That considered, the uncorrected coefficients of these tests can still be used to broadly assess differences between the distributions.

Table 7.6 shows a single slice of this coefficient matrix corresponding to the behavioural measure of *duration* for each application-game combination described in Figure 7.5. From this slice alone it is clear that the coin-flip game on the fck.com application stands out against almost all others in terms of the size of the K-S II coefficient. The K-S II coefficient can be interpreted as the probability that the two distributions are different. It is therefore possible that if the players of the other games are human, then fck.com coin-flip players are not. The results of these tests across each of the behavioural measures in the matrix appears to indicate non-zero differences between the fck.com coin-flip players against players of all other provider-game combinations. Add to this that the fck.com coin-flip game amassed

Data: Behavioural measures for all players

Result: K-S II tests between measures for each application-game combination

```

data = [player measures for each app-game combination];
allMeasureTests = [];
for measure in measures do
    // 2D matrix for one measure;
    testResults = [];
    for column in data do
        for row in data do
            | testResults.append(KStest(column, row));
        end
    end
    allMeasureTests.append(testResults);
end

```

Algorithm 1: Two sample Kolmogorov-Smirnov tests for each behavioural measure between groups of players, where each group represents the players of a single game on a single application. The Python implementation used in this study is available as part of the *gamba* library at www.gamba.dev.

13,877 unique players over its lifespan of 209 days compared to 567, 293, and 396 players among its other three games, it appears unlikely that the majority of transactions to this game are human in origin.

The two sample Kolmogorov-Smirnov test results between the *dice2win* coin flip players and the *fck.com* double dice players are also higher than any other non *fck.com* pair. Yet with no other pairs indicating distributional differences with this group this may be an artefact of the choice of game, or may be coincidental given the number of tests conducted. In each case, this anomaly invites further exploration but is considered out of scope of the present study.

Under the assumption that each of the remaining provider-game pairs' transactions originate from human players - which I found no evidence to refute - the *fck.com* coin-flip transactions were discarded. This left 8 application-game combinations of interest, whose 10,357 unique players' behavioural measures - using 1,743,478 transactions - were combined into a single data set, as performed in existing work in gambling behaviour analysis [2]. A graphical breakdown of these application-game combinations is provided in Figure 7.5. As with the matched transactions described above,

Table 7.6: Two sample Kolmogorov-Smirnov test results for player durations across all provider-game combinations. † denotes a significant result ($p < 0.01$) and coefficients greater than 0.35 are highlighted. Key: d2w = dice2.win, fck = FCK.com, eroll = Etheroll.com, cf = coin flip, sd = single dice roll, dd = double dice roll, oh = 1–100 roll.

DApp		d2w				fck			eroll	
	Game	cf	sd	dd	oh	cf	sd	dd	oh	oh
d2w	cf	-								
	sd	0.16	-							
	dd	0.24 [†]	0.10	-						
	oh	0.17	0.03	0.08	-					
fck	cf	0.55[†]	0.39[†]	0.46[†]	0.39[†]	-				
	sd	0.22 [†]	0.08 [†]	0.05	0.05	0.43[†]	-			
	dd	0.40[†]	0.25 [†]	0.16 [†]	0.23 [†]	0.59[†]	0.20 [†]	-		
	oh	0.34 [†]	0.19 [†]	0.10 [†]	0.17 [†]	0.54[†]	0.13 [†]	0.07	-	
eroll	oh	0.09	0.12	0.19 [†]	0.11	0.49[†]	0.14 [†]	0.34 [†]	0.27 [†]	-

the table of behavioural measures calculated for each unique address in this study is available through <https://osf.io/8bfyj/>.

Cryptocurrency gambling behaviours

Table 7.7 presents the behavioural measures described in the Measures section above for the cohort of players in the remaining transaction set. The majority of the measures have heavily skewed distributions, which limits the descriptive power of the parametric statistics presented. This table therefore extends LaBrie et al’s original metrics [2] by including the inter-quartile ranges of each of the measures, plus the coefficients of a one sample Kolmogorov-Smirnov test for normality as reported by LaPlante et al [100].

The results show that with a median duration of 1 day and frequency of 100%, the typical player of decentralised gambling applications bets in a non-committal, and non-intense way. This contrasts LaBrie et al’s original findings on regular casino players, who with a median duration of 246 days and frequency of 7% bet across a much longer term. This contrast may be explained in part by the youth of the applications studied here. Add to this a median bet count of 11 and we may assume that this typical player would play for one short session on a single application and then cease play altogether or move to another application. This considered, the inter-quartile range for the duration indicates a portion of players remaining engaged for over a week of play. Combine this with the inter-quartile ranges for both

Table 7.7: Gambling behaviour of 10,357 decentralised gambling application players including a one sample Kolmogorov-Smirnov (K-S) test for normality. All K-S test statistic values are significant at the $p < 0.01$ level, STD = standard deviation, IQR = inter-quartile range.

Metric	Mean	STD	Median	IQR	K-S
Duration (days)	30	81	1	10	0.841
Frequency (%)	76	36	100	50	0.966
Number of Bets	168	992	11	62	0.841
Mean Bets/Day	23	48	6	21	0.841
Mean Bet Size	1.15	11.8	0.11	0.35	0.504
Total Wagered	213.77	2451.85	1.40	16.59	0.504
Net Loss	2.91	49.86	0.04	0.71	0.213
Percent Loss	10.9	112.1	5.3	52	0.548

the frequency and number of bets measures and we observe a wide range of possible behaviours between the 25th and 75th percentiles of the sample, across the measures calculated. This sentiment is shared in the number of bets placed per betting day, which, with a median of 6 and IQR of 21, encapsulates a wide range of possible behaviours for the majority of the sample.

The top four behavioural measures also present the highest one sample K-S test statistics of all of the measures. This is most likely an artefact of the heavily skewed nature of these measures, with thorough investigations of outliers across each measure representing an interesting area of future work.

The financially oriented measures, including the ETH per bet and the total amount wagered show similar oddities to the results regarding duration and frequency. With a median bet size in ETH of 0.11 (approximately equivalent to \$14, see <https://www.coinbase.com/price/ethereum> for exchange rate data used throughout this work) and total amount wagered of 1.40 ETH (approximately \$200), the typical player's spending is high considering the short duration of play. The granular and longitudinal nature of the transaction data prepared as part of this work mean that questions surrounding this behaviour, and its relation to external harm related variables, can be explored in greater detail in further work, but are not expanded upon here.

The most comparable measure presented here with other gambling activities is the net and percentage loss measures, which with median values of 0.04 (ETH) and 5.3% respectively indicate modest losses for the typical

Table 7.8: Non-parametric Spearman rank-order correlations between all behavioural measures for decentralised gambling application players. All values are significant at the $p < 0.01$ level. Coefficients of magnitude greater than 0.70 are highlighted.

Measure	Duration	Frequency	# Bets	Bets /day	Eth /bet	Total Wagered	Net loss	% loss
Duration	-							
Frequency	-0.89	-						
# Bets	0.63	-0.45	-					
Bets/day	0.35	-0.19	0.93	-				
Eth/bet	0.16	-0.10	0.26	0.24	-			
Total Wagered	0.53	-0.39	0.84	0.78	0.72	-		
Net loss	0.12	-0.10	0.15	0.14	0.15	0.20	-	
% loss	-0.10	0.06	-0.15	-0.12	-0.07	-0.14	0.67	-

player. As with other financially oriented measures, when framed in terms of the median duration this equates to a loss of 5.3% of the total amount bet per day.

Unlike the top four measures presented, the financially oriented measures do not present such high K-S test statistics, so are likely drawn from less extreme distributions. That considered, with test statistics of 0.504, the ETH per bet and total amount wagered measures still cannot be effectively described using parametric methods. As such, the means and standard deviations for each of the measures are reported in line with existing literature, but in this domain do little to develop our understanding of typical transactional behaviour.

Relationships between behaviours

As with previous work exploring the behavioural measures used in this work [61][100], heavily non-normal distributions mean that rank-order correlations are preferred over their parametric equivalents. Table 7.8 presents Spearman rank-order correlation coefficients between all of the behavioural measures calculated for players of all games combined, excluding the fck.com application's coin-flip players.

Of particular interest in these coefficients are those of substantial magnitude, as highlighted. I find that, as expected, frequency is negatively correlated with duration - this makes sense as given a larger number of possible days on which to place a bet, the probability of a player not placing

one on a given day naturally increases. The measure of duration does not appear substantially correlated with any remaining measures, with moderate values for both number of bets and total amount wagered. These each loosely support notions that the longer an individual uses a decentralised gambling application, the more bets they will place and the greater their total amount wagered will become. These each also make logical sense in the context of the gambling games these applications present.

Apart from its correlation with duration, the measure of frequency does not appear to relate to any other measures in any substantial way. With a coefficient of 0.46, its correlation with the number of bets an individual makes also makes intuitive sense. The more frequently a player places bets, the more bets they are likely to place over their gameplay career.

The number of bets appears strongly correlated to both the number of bets per day and the total amount wagered for users of decentralised gambling applications. With a coefficient of 0.93 - the strongest of all pairs - it is clear that the number of bets an individual places over their duration of play directly relates to the number of bets they are likely to place on a given day. The number of bets measure also relates strongly (0.84) to the total amount wagered. Unsurprisingly, the number of bets an individual places on a given day is also strongly correlated (0.78) with their total amount wagered. As with other relationships between measures, this makes intuitive sense in the context of gambling games but nonetheless contributes to establishing a baseline for human players of such games.

The final coefficient of interest, and that of most potential scientific significance, is that between the ETH per bet and the total amount wagered. With a reported coefficient of 0.78, these results suggest that those who place larger bets are more likely to wager larger total amounts over the duration of their betting careers. The implications of this finding are deferred to the discussion. However, this appears to suggest that this measure may be an important predictive indicator in the cryptocurrency domain. It may assist in terms of identifying the potential for financial harm via unsustainable spending among players - a finding in line with existing work in player behaviour tracking research [1].

Both the measures of net loss and percent loss do not appear meaningful in relation to the other measures reported in this work, so will not be discussed in detail. One can now move on to report descriptive statistics regarding

Table 7.9: Non-parametric descriptive statistics of the behavioural measures for the top 5% most heavily involved bettors by total amount wagered, and the other 95% of players. All one sample K-S test statistic values are significant at the $p < 0.01$ level indicating the data for each measure is non-normally distributed.

Measure	Top 5% ($n = 518$)			Other 95% ($n = 9,839$)		
	Median	IQR	K-S	Median	IQR	K-S
Duration (in days)	35	120	0.91	1	7	0.84
Frequency	50	78	0.98	100	50	0.97
Number of bets	644	1660	1	9	47	0.84
Bets per day	68	77	1	5	18	0.84
ETH per bet	1.84	5.61	0.53	0.10	0.28	0.50
Total wagered	986.39	1759.01	1	1.10	10.89	0.50
Net loss	10.3	102.6	0.56	0.04	0.6	0.22
Percent loss	0.9	7.6	0.38	6.6	57.6	0.56

the most heavily involved bettors in the data set, and contrast them to the majority of low and moderately involved bettors.

Heavily involved bettors

Heavy involvement by any of the behavioural measures used here may be detrimental to the individuals affected. For example, those most heavily involved in terms of the duration of their play will naturally have less time for other commitments, or those with large net losses may face financial repercussions should their income not support such expenditure. I explore heavy involvement with respect to total wagered, as it has the most obvious financial repercussions for the individuals in the cohort. This follows LaBrie *et al.*'s rationale for exploring the same measure in a cohort of casino gamblers [2]. I include LaBrie *et al.*'s original figures for quick comparison in Table 7.10, although such comparisons are heavily nuanced given the differences between decentralised and regular online casinos.

Table 7.9 presents each of the descriptive statistics for each of the behavioural measures, for both the top 5% most heavily involved bettors by total wagered, and the remaining 95% of the sample. Parametric statistics for both cohorts are not reported given their heavily skewed nature as described in the previous section.

These results begin with substantial differences in the typical duration of play between those most heavily involved and the remaining 95% of the

Table 7.10: Gambling behaviour of extreme 5 and 95% subgroups of casino bettors, reprinted from LaBrie et al's 2008 study [2].

Measure	Most involved casino bettors top 5% ($n = 212$)		Other 95% of participants ($n = 4,010$)	
	Mean (SD)	Median	Mean (SD)	Median
Duration	476 (232)	529	290 (233)	246
Frequency	24% (17)	20%	16% (21)	7%
Number of bets	24,558 (36,779)	10,465	2,403 (7,819)	486
Bets per day	285 (344)	188	107 (176)	46
Euros per Bet	213 (682)	25	25 (97)	4
Total Wagered	345,579 (354,890)	233,195	10,338 (19,360)	2,284
Net Loss	8,746 (11,213)	6,698	422 (939)	107
Percent Loss	2.6 (3)	2.5	8.0 (12)	5.9

sample. Whilst the typical player in the majority only plays for a single day, placing approximately 9 bets in total, the typical heavily involved bettor plays for over one month, placing over 600 bets. Furthermore, the typical heavily involved bettor appears to spread these bets over the month, betting approximately every other day, just under 70 times. The difference between the typical bets per day multiplied by the typical number of betting days per month ($70 \times 15 = 1050$), and the typical number of bets alone (644), indicates a difference in the range of behaviours these players are exhibiting. Exploring these differences represents a key area of future work.

Each of the bets of the typical heavily involved bettor are also not insignificant in size, being almost 20 times higher than the typical player in the majority of the sample - a median 1.84 ETH (roughly \$200) compared to 0.1 ETH (roughly \$10). The most dramatic difference, and most concerning for the players affected, is the difference between the median total amount wagered between the most heavily involved bettors and the remaining players. With a median of almost 1,000 ETH (equivalent to approximately \$100,000), it dwarfs the median 1.1 ETH (\$110) presented by the majority of bettors. This proportional difference is consistent with LaBrie et al's original study on regular online casino gamblers (see Table 7.7), but appears to amplified in decentralised gambling application use. This difference of an almost $1000\times$ greater total amount typically wagered by heavily involved players compared to the majority of players is a key finding of this work.

As with the behavioural measures reported for the entire sample in Table 7.7, the inter-quartile ranges of each of the measures leaves a wide range of

potential transaction behaviours for those in the top 5%. This includes the duration measure, with players engaging with the decentralised gambling applications across a range of 0 to 155 days or more. This holds for the frequencies, with some heavily involved players betting every day throughout the duration of use, and some betting only a few times with large wagers. Most varied in terms of non-financially oriented measures is the number of bets placed, which presents an inter-quartile range of over 1,600 for the top 5% compared to 47 for the majority. This is of particular interest regarding the use of this data for transaction pattern analysis, a potentially fruitful area of research extending this work, and discussed in more detail below. With so comparatively few transactions made by the majority of players, further studies using this data should use behavioural measures which account for this difference.

Other widely varying measures include the total amounts wagered and the net loss. The median values of total amount wagered are 986 and 1.1 with inter-quartile ranges of 1759 and 11 respectively between cohorts. Net loss shows similar ranges with medians of 10.3 and 0.04 with inter-quartile ranges of 103 and 58 respectively. This develops the previous finding that among the top 5% of most heavily involved players, a wide range of potential patterns exist, confirming the existing idea that there is no single behaviour indicative of heavy involvement, rather a spectrum of potential patterns and behaviours which each result in large total expenditures.

Lastly, with respect to the descriptive statistics presented, the percent loss between the most heavily involved players and the majority presents a counter-intuitive result. With a larger total amount wagered, the losses one may anticipate for the typical heavily involved individual would be high, although, in decentralised gambling applications it appears to be the opposite. With a median percent loss of just 0.9 and an inter-quartile range of 7.6, the typical heavily involved bettor does not appear to lose the amount they wager in as varied a fashion as the other 95% of the sample. These values align with Labrie et al's original work, which also reports lower percent losses for heavily involved bettors (2.5%) than for the majority of the population (5.9%). This may be an artefact of the provable fairness of these games as described in the Data Sample section above, where players can be certain of the amount the 'house' is taking from each bet, or it may be a result of extensive repeated play, where the range of potential losses is effectively

smoothed by the larger sample available for each player. In the case of the majority of players, a median percent loss of 6.6 and inter-quartile range of 57.6, suggests large relative wins and losses for the relatively small bets they place. This finding differs from the original work, but makes logical sense given the non-committal and non-intense behaviours described above for the typical player of decentralised gambling applications.

The one sample K-S statistics reported for the behavioural measures of the heavily involved portion of players and the remaining 95% indicate several measures of interest for future work. Specifically, the differences between the first four measures (duration, frequency, bet count, and bets per day) do not appear substantially different from one another. The differences in distributions between the total amounts wagered however are vastly different, with a coefficient of 1.00 (to 2 decimal places) for heavily involved bettors compared to 0.50 for the majority. This may be a fruitful area of further exploration, as the underlying distributional differences for these measures may be used in conjunction with other measures to predict heavy involvement.

7.3.4 Discussion

This study presents the first ever analysis of decentralised gambling transactions on the Ethereum blockchain. Decentralised gambling, and the contract components of their architectures, present significant regulatory challenges [5], whilst simultaneously offering rich transaction level data for research. Whilst this transaction level data exists in large quantities, we have shown that the entire set is not immediately useful for research given the likely presence of non-human players. This means that although a large, publicly available, *in-vivo* data source for player behaviour research has emerged, scholars must take care when using it to solve existing problems, especially when exploring issues around disordered transaction patterns and player behaviour clustering.

Non-human players in decentralised gambling applications

The first distinct analysis involved employing statistical tests to detect differences between transactions of (likely) human and non-human origin. To this end, I found that performing two sample Kolmogorov-Smirnov tests between behavioural measures, and between games provided by decentralised

gambling applications, can be effective for detecting the presence of players whose transactions stand out against those in other games. This simple method invites improvements, but shows that relying on distributional differences between human and non-human players is enough for meaningful distinction at this early stage.

Importantly, I hold the assumption that of the nine application-game combinations we explored, the one that stands out as different is not being transacted with by human players, as opposed to the other way around. Under this assumption, one may suggest that the reason it differs so substantially from others is that the majority of the players are in fact not human, but instead are cryptocurrency spending/betting algorithms designed to transact with the application, potentially to inflate perceived popularity. Exploring motivations behind deploying algorithmic interactions with these applications presents an interesting but tangentially related area of future work.

Cryptocurrency gamblers and behavioural relationships

The second and third analyses described in Section 7.3.1 above aimed to describe the gambling behaviours of users of decentralised gambling applications, and assess the relationships between these measures. These results suggest, as with similar existing work [2], that the distributions of all behavioural measures are significantly skewed, and therefore benefit from the application of non-parametric statistics. Applying such statistics, and without breaking down the sample of players into meaningful sub-samples, the typical user of decentralised gambling applications does not appear to be heavily involved, and does not appear to place a substantial number of high of bets. However, this description fails to capture the most important aspect of the findings in this study, which are that those most heavily involved in the use of decentralised gambling applications appear to spend significantly more than both the majority of the population, and more than heavily involved gamblers in other types of online gambling. Exploring this relationship further and breaking down differences in terms of the behavioural measures calculated for each player, presents a fruitful area of further work if findings building on previous studies are to be translated to this new domain.

Furthermore, this study's design draws heavy inspiration from early work describing online casino game players. The data available to the original researchers took a daily aggregate form. This means that the

behavioural measures they devised to describe cohorts of players perhaps do not capture the depth of insight available when using individual transaction level data as available via cryptocurrency transactions. There may therefore be behavioural measures which appear inaccessible at the daily aggregate level, such as average gambling session length or average rate of spending. To the author's knowledge, studies in the field of player behaviour tracking have not yet explored such granular measures, nor applied them to data sets across different types of online gambling. That considered, measure-oriented work such as Kainulainen's [110], which describes a new measure of risk taking specific to gambling, presents the opportunity to apply new techniques to gain deeper insight on player behaviours.

Heavily involved cohort characteristics

The final analysis aimed to provide an epidemiological description of the gambling behaviours of an empirically determined group of heavily involved gamblers. The results regarding this cohort of players, identified as heavily involved by total amount wagered, suggest a number of important discoveries. Firstly, although the typical heavily involved player spends the equivalent of over \$120,000 during a 35 day period, the losses they typically incur as a percentage of their amount wagered are under 1%. This means that although their expenses dwarf the majority of players by over 1000×, they do not appear to be losing as much proportionally as the majority of players, who, when placing approximately \$105 worth of bets in total over a one day period typically lose under 6%, or \$7. It is important to note that this difference in losses between heavily and non-heavily involved players is not unique to decentralised gambling applications, as evidenced by LaBrie et al's original findings (Table 7.10).

Another important result of the analysis regarding heavily involved bettors is the typical difference in bet size, with heavily involved players wagering just under 20× more than their low to moderate counterparts. This result can be used to inform further research on the use of cryptocurrencies for gambling, and the analysis of their transactions, for the early detection of unsustainable spending, for example. This is just one of many possible - and much needed - avenues of work extending these findings into the domain of responsible gambling analytics.

Limitations

The analyses performed here are subject to many of the same limitations of the use of online gambling data for behaviour tracking research generally [1]. These include issues surrounding the generalisability of findings. In the context of the use of cryptocurrencies for gambling - specifically through decentralised gambling applications - is unclear whether the analysis undertaken here will have similar results across other comparable applications. Furthermore, it may be the case that the behavioural patterns uncovered here are incomplete as true player gambling behaviour may be spread across several unobserved applications in addition to the applications discussed here.

Such fundamental limitations can not be completely negated through experimental design. However, future work should focus on increasing the sample size, both to more applications and more players, which may address the issue of generalisability regarding decentralised gambling applications.

An additional point may be made regarding the transaction matching process performed which pairs incoming and outgoing transactions. The data that this analysis was conducted over involves a complete record of each player's ingoing and outgoing transactions. However, it does not contain a reliable temporal ordering for this data. In order to create a more useful data set than the incoming and outgoing transactions in isolation, they can be matched such that an outgoing transaction chronologically following an incoming transaction from the same address can be taken to be the payout of a previously placed bet, but that other candidate transactions may be considered in the case that one transaction is completed ahead of another.

Perfect transaction matching, as described above, is unnecessary for the methodology used in this paper, as all behavioural measures computed use the aggregation of an individual's ingoing and outgoing transactions. For example, the behavioural measure of percentage loss for a given player only requires the sum of their bets and the sum of their payouts. However, one might imagine the calculation of more sophisticated behavioural measures that do require matched data in order that more sophisticated analyses might take place. For example, one might attempt to calculate the phenomenon of 'chasing losses' by measuring the extent to which players place larger bets after losing money on a prior bet. Such an analysis is not possible using the data set outlined above, as any given payout could not be conclusively

matched to a single bet. This matching process is briefly mentioned here as it will be essential for future work in this area at the individual bet and risk analysis levels - both techniques are considered out of scope of the present study.

Other limitations relate to the nature of the applications themselves in comparison to other online gambling platforms. Specifically, each of the applications used here - and all decentralised applications atop cryptocurrency networks - must use cryptocurrencies or similar tokens by design. This means that although the real world value (e.g. in US\$ or GBP) for any amount of cryptocurrency can be determined in real terms, it is unclear whether or not this relationship affects wagering, and in what way. An area of future work exploring this relationship may investigate the distributions of bet sizes, and may uncover more detailed findings in terms of how decentralised gambling differs from other online gambling. These studies may also help in understanding how the use of virtual goods and currencies affects the behaviour of players with respect to spending. In this vein, comparisons with other uses of cryptocurrency technology, such as the development of crypto-games [9], may provide a useful basis for comparison.

The differences between these applications and other online gambling providers also inherently affects the populations who use them. This means that the sample of players considered in this work is a sub-population of individuals who have purchased cryptocurrencies - a volatile [9] and technologically sophisticated means of facilitating e-commerce [5].

A final limitation of this work, given the context of recent advances in player behaviour tracking research, is that it only explores simple behavioural measures based on those used to explore casino gamblers [2]. It therefore does not reach into more advanced analytical methods for describing, classifying, or predicting player behaviours. This includes work by Fiedler which explores more granular behavioural measures [63, 65], multiple studies by Percy [72] and Dragičević *et al.* [102], which employ neural networks and other machine learning methods for responsible gambling, and other data mining procedures for identifying high risk gamblers as done by Philander [81]. In order to apply supervised machine learning as in these studies, labelling heuristics for players should also be explored.

Future Work

The analysis presented here recreates that of a series of papers originating from Harvard Medical School [61, 2, 100]. Since that series was first published, a number of other descriptive measures have been used such as the intensity, variability, frequency, and trajectory of a player's bets [3], and more specific variables such as the number of betting sessions and total time spent betting [63]. Extending the present study by exploring player behaviours across these dimensions would give a more complete picture of the player base of decentralised gambling applications, and would give stronger grounds on which these transactions may be compared with other types of gambling.

A second avenue of research extending the descriptive and test statistics reported here is the use of this data for identifying and predicting high risk gambling. Existing work has identified transaction patterns and behaviours to be markers of high-risk play [69, 106] - exploring such methods in this new domain may therefore help identify those at risk, and better describe the way these applications are used. The development of such identification methods may spur regulators and policy makers to further explore cryptocurrency exchanges, whose operations provide financial access to these applications. An obvious and useful first step would be formally requiring currency exchange reporting for responsible gambling analysis.

Finally, the findings presented here may be tentatively mapped to other forms of gambling in which similar work has been reported. Generalisations drawn from such mappings may require further data gathering from both additional cryptocurrencies, such as the EOS network, and more applications on the Ethereum network as described in this study and elsewhere. Increasing the sample size of players, both human and otherwise, represents a strong second step in creating reliable and generalisable findings, which extend this work.

7.3.5 Conclusion

In this study, 2, 232, 741 transactions to and from three decentralised gambling applications, originating from 24, 234 unique cryptocurrency addresses, were gathered, and four distinct analyses performed. These findings suggest that not all transactions to decentralised gambling applications originate from human players, making data cleaning crucial in all further academic work

concerning this type of data. The analysis presented found a pairwise two sample Kolmogorov-Smirnov test across players behavioural measures to be effective in distinguishing assumed non-human players. Of transactions believed to originate from human players, the behavioural measures computed naively were found to describe non-intensive but moderate spending over a short duration for the typical player. This description was then found to mask a small portion of heavily involved bettors, whose typical bet size appears to be almost $20\times$ larger than their non-heavily involved counterparts, and their total amount wagered appears to be over $1000\times$ larger over their duration of play. Contributions presented in this paper are two-fold; the work presented primarily illustrates the power and scale of transaction data that decentralised gambling applications can provide gambling researchers. Secondly, it describes a large cohort of players from three such applications, and uncovers extreme behaviours, such as large bet sizes and substantially larger total wagering among heavily involved players. This work should draw attention to cryptocurrency transactions as a tool for large scale *in-vivo* gambling research, and presents a robust foundation upon which multiple avenues of further analyses can be pursued.

7.4 Summary

This chapter has presented the first ever study of player behaviours in decentralised gambling applications using actual cryptocurrency gambling transaction data. This data was taken from the market leaders established in Study 2 presented in Chapter 6, and extracted from the Ethereum blockchain by applying knowledge of the application architectures from each of the case studies. This study found that in comparison to existing epidemiological profiles, gamblers in the cryptocurrency domain engage less overall, but the most involved portion of the population engages significantly more than their non-heavily involved counterparts (Table 7.9). Additionally, a subset of suspected non-human players were identified by measuring the difference in distributions of behavioural measures - a technique not yet applied in other player behaviour tracking studies, but which successfully detected an entire game (FCK coin flips) which contains transactions which are likely non-human in origin. This study provides an important epidemiological profile of players in this emerging domain, and can be used to frame more

detailed behavioural profiling methods as presented in Chapter 9. The next chapter in this thesis branches into the blockchain games domain, extending the profiles established in this chapter to this adjacent but related domain.

Chapter 8

Inside Blockchain Games

“Once more unto the breach, dear friends, once more”

William Shakespeare

Some mechanisms in digital games exhibit mechanical similarities to gambling [126]. One such type of mechanism is commonly referred to as loot boxes [127][128] or randomised reward mechanisms [4], and are of academic and phenomenological interest to player behaviour tracking researchers given their potential for gambling related harms [128]. Blockchain games can be considered a type of digital game, and therefore may contain randomised reward mechanisms. In existing literature, no studies have explored the behavioural profiles of players interacting with such mechanisms in particular, specifically through the lens of player behaviour tracking research [1]. This chapter addresses this by providing a population level behavioural analysis of players engaging with randomised reward mechanisms.

This chapter, like the previous chapter, aims to apply all of the knowledge from Chapters 2, 3, and 5, to cryptocurrency transaction data - this time to the domain of decentralised gaming applications, or blockchain games. Following the introductory Section 8.1, Section 8.2 discusses the architecture of a leading blockchain game (CryptoKitties), the understanding of which can then be used to generate meaningful data sets from the raw transaction data available on the blockchain. Section 8.3 then presents the fifth major study in this thesis which applies the same behavioural measures as used in Study 4, this time to a set of randomised reward mechanism function calls in the CryptoKitties application. The application of behavioural measures in this

way allows a direct comparison between behaviours in these studies, enabling a broader understanding of similarities and differences between these two distinct but related application types. Finally, Section 8.4 concludes this chapter, summarising the key findings from this study.

8.1 Introduction

So far, this thesis has explored the use of cryptocurrency technology in applications marketed explicitly for gambling. Before moving to more detailed behavioural analyses in the gambling domain, a brief detour into the closely related realm of crypto-gaming is presented. These so-called blockchain games share many of the mechanical properties of the decentralised gambling applications analysed in the previous chapter, but also incorporate gameplay elements similar to many digital games, such as quests, collectable items, and so on. As they contain randomised reward function calls [4], which act in much the same way as bet placement functions described in previous chapters, an analysis of the population of players through the lens of player tracking is warranted. This chapter therefore presents two studies, the first maps the rate of adoption of this new domain, providing broader context for the second, which establishes a behavioural baseline for comparison with the gambling applications in the previous chapter.

Smart contracts can be used in digital games to store in-game tokens or currencies on a blockchain. Once stored in this way, these tokens can be transacted in the same way as the underlying cryptocurrency. This gives players the ability to trade their virtual goods, and enables interoperability of virtual goods between games. This is notably different to centralised games in which a single server or cloud holds all player data and falls under total governance of the game developer/publisher. In the latter, interoperability and trading is only possible if the centralised developer allows it, whereas in the former, players have full control over their virtual goods and how they are used. Crucially, all these transactions are stored atop a cryptocurrency blockchain, so can be directly accessed and decoded in order to gain insights into the behaviours of players at the population and individual levels.




Figure 8.1: The general form of randomised reward mechanisms described by Nielsen and Grabarczyk [4].

8.2 Application Architectures

Before applying more granular behavioural measures to blockchain games transaction data, a brief analysis of application architectures is presented. Unlike the decentralised gambling applications discussed in the previous chapter, blockchain games can exhibit vastly different architectures given their practically limitless creative directions and resulting programmatic requirements. This diversity becomes particularly challenging when deeper explorations of player behaviours - which require decoded transaction data - are needed, as there is an upper limit to the depth of understanding a lone or even team of researchers can build across a sample of this size. For this reason, this thesis presents a case study of the popular blockchain games CryptoKitties, following a more concrete discussion of which function calls are of academic interest in the context of gambling.

8.2.1 Randomised Reward Function Calls

Randomised reward mechanisms as described by Nielsen and Grabarczyk [4] consist of a structure with a very general three part form, reproduced in Figure 8.1. These three parts - some eligibility criteria followed by a random procedure followed by a reward - can be identified in the internal mechanisms of crypto games via examination of the smart contracts associated with these games. For example, CryptoKitties is a collectable kitten breeding game which uses smart contracts to process payments and calculate outcomes. The CryptoKitties' smart contracts are *mostly*¹ open source and can be inspected via the Etherscan platform as described in Section 3.4. In the case of CryptoKitties, a `breedWithAuto` function exists whose purpose is to commit two kittens to a breeding process such that after some time delay a new kitten will be created. This new kitten creation process hinges on a secret

¹The genetic mixing algorithm used by the CryptoKitties smart contract is not publicly available in raw source code form, only in the encoded form stored on the Ethereum blockchain.



What is CryptoKitties?

CryptoKitties is a game centered around breedable, collectible, and oh-so-adorable creatures we call CryptoKitties! Each cat is one-of-a-kind and 100% owned by you; it cannot be replicated, taken away, or destroyed.

Figure 8.2: Description of CryptoKitties, taken from <https://www.cryptokitties.co/>, accessed 06/07/2021.

probabilistic gene mixing function, which determines the characteristics (and therefore the value) of the resulting virtual kitten, which is then returned to the original caller of the function. This relatively simple mechanism matches the general form of a randomised reward mechanism in [4], and is therefore of interest in subsequent transaction analytics as part of this thesis, as transactions representing function calls to this smart contract directly capture player's interactions with this randomised reward mechanism.

8.2.2 Case Study: CryptoKitties

Study 3 in Chapter 6 found that CryptoKitties is a market leader in both user count, transaction count, and transaction value. CryptoKitties is a collectable kitten breeding game created by DapperLabs in 2017 [40]. The objective of this game is to repeatedly breed virtual kittens in order to create rarer, more valuable kittens which can then be sold or bred again (See Figure 8.2). Each kitten exists as an ERC-721 token, which means they are non fungible and can be used in other interoperable applications by their owners once created or purchased. This collection of compatible applications is referred to as the Kittyverse by DapperLabs, but is not the focus of this case study. Instead, this case study focuses on the core functionality of the blockchain game, as described in the application's 'core' contract.

Contract ABI

Given the CryptoKitties application's richer functionality than the decentralised gambling applications analysed in the previous chapter, its smart contract is naturally more complex. Specifically, the application has a 'core' contract, which contains all of the core game logic for breeding and transferring the kittens, but also has a number of other contracts which contain sales auction, siring auction, and other broader functionality. This modular design makes research into specific user actions within the application relatively simple, although an understanding of how those user actions are represented programmatically is still required. This case study therefore focuses explicitly on the KittyCore smart contract currently in operation at `0x06012c8cf97bead5deae237070f9587f8e7a266d`.

This contract contains 64 functions (compared to a maximum of 38 (Etheroll) in the previous case studies), which range in functionality from assigning new executive addresses (`setCEO`, `setCFO`, etc), through to game management (`setGeneScienceAddress`, `setAutoBirthFee`, `ContractUpgrade`), and basic Kitten functionality (`createPromoKitty`, `breedWithAuto`, `createSiringAuction`) to name a few. Many of these functions will never be interacted with by players, but exist instead as the blockchain equivalent of management commands which can only be executed by those with authority within the application. Such functions are therefore not of interest for understanding player behaviours, so are not discussed further. Additionally, several functions do not correspond to direct player actions, but are instead subroutines within a player action. For example, the `isReadyToBreed` function checks whether a supplied Kitten can engage in the breeding process. This function is only called in the `breedWithAuto`, `createSiringAuction`, `bidOnSiringAuction` functions, and would therefore never be called directly by a player, but is instead called when they call one of these three functions. The existence of subroutines highlights the importance of developing a deep understanding of a given application before selecting function calls for analysis. As in the gambling domain, this is a resource and time intensive process which can become prohibitively costly for all but the simplest of blockchain games.

As described in Section 8.3.1 above, the CryptoKitties `breedWithAuto` function captures the interaction of players with the probabilistic breeding process which meets Nielsen and Grabarczyk's definition of a randomised

reward mechanism [4]. For this reason, only an understanding of this function is required for further analysis in this domain. Additionally, unlike the applications in the decentralised gambling domain, there is no question of which specific game is being played, or which parameters are of interest, as it is engagement with the function itself which is of interest. This means that an audit of the parameters and internal operations within this function are not required in order to generate meaningful data sets for analysis using the raw blockchain data available on the Ethereum network.

8.3 Study 5: Behavioural Distributions in Blockchain Games

The fifth study in connection with this thesis concerns understanding the behavioural profiles of blockchain games players when interacting with randomised reward mechanisms, which in the cryptocurrency domain (and therefore in this thesis) are considered a unique form of decentralised gambling application. This study specifically aims to establish an epidemiological baseline of blockchain games players using the same behavioural measures applied in comparable gambling research. This can be achieved by isolating randomised reward function calls in these games, as these are the most mechanically similar user actions to the placement of bets in casino games [4]. The subsequent behavioural measures therefore provide an empirical description not of how users engage with blockchain games in general, but how they engage with the most gambling-like features within these games specifically. This study therefore advances our understanding of blockchain games in two key ways. The first by providing an epidemiological account of engagement with randomised reward mechanisms in the blockchain games domain, and the second by providing an empirically based comparison of this engagement with actual gambling also in the cryptocurrency domain.

8.3.1 Introduction

The previous study in this thesis has shown that the use of Ethereum blockchain games has shown consistent growth despite the volatility of the underlying cryptocurrency. Should this rate of growth continue, ever more players will choose to play blockchain games, and existing centralised game

developers may choose to integrate cryptocurrency technology into their existing intellectual properties. This potential for increased adoption of this new technology, and the almost 700,000 existing unique accounts, means understanding how users are interacting with these games is essential to fostering informed academic and regulatory discussions around their use. This understanding is made even more crucial in the specific context of problematic player spending and monetisation - a topic within games research and human-computer interaction aimed at minimising harm in digital games in which there is much debate [127][128][129]. Importantly, evidence within this debate is lacking in the blockchain games domain, with only a few studies exploring their mechanisms [40], and fewer still applying behavioural data analyses.

Like their centralised counterparts, blockchain games employ a number of monetisation techniques, the most obvious of which being the intergration of randomised reward mechanisms (RRMs) into the core gameplay loop (See Section above for a description of RRM). However, unlike their centralised counterparts, transactions representing player interactions with these mechanisms are stored in the publicly available Ethereum blockchain, rather than on private servers. This means they can be accessed through blockchain explorers such as Etherscan, or through a fully synchronised archival Ethereum node, by any researcher with an internet connection, and a detailed knowledge of the architectures of such applications. This full historical data availability represents a paradigm shift in digital games research, and enables a host of novel research directions, the most relevant to this study being the use of these transactions for behavioural profiling of players through the lens of player behaviour tracking.

Whilst studies of CryptoKitties transaction data have been published in recent years [130], none have taken a player behaviour tracking oriented approach, making this study the first to allow direct comparisons with gambling research. This is important as it adds ongoing discussion around the similarities between randomised reward mechanisms and gambling, which if mechanically similar, should yield similar behavioural profiles.

Randomised Reward Mechanism

The CryptoKitties game includes several randomised reward mechanisms which are responsible for initiating a breeding process and generating new

Kittens. One such mechanism for automatically breeding Kittens is isolated in the smart contract as a `breedWithAuto` function, accepting a fixed amount of ETH as payment and returning a new Kitten to the cryptocurrency address which invoked the function call. This function's architecture follows Nielsen and Grabarczyk's generalised form of randomised reward mechanisms [4], and has obvious similarities to gambling via simple casino games, although here the 'payout' in the casino sense is the new kitten itself rather than return in the same currency the payment was made (ETH). Importantly, this architecture means that by decoding Ethereum blockchain transactions one can identify those which invoke this function and group them into a single data set which represents the entire CryptoKitties player base' interaction with this mechanic. One cannot, however, easily value the resulting Kittens, making analysis in the loss domain (See Chapter 5) difficult - this limitation is discussed as encountered.

8.3.2 Method

The application of behavioural measures in the context of digital games includes those in the time and cost domains described in Chapter 5. This is because, as mentioned above, the payout takes the form of a token (kitten) rather than an amount of currency. This means that to properly value the payout, one would need to use third party data from an NFT exchange and pair it to the transactions. The global nature of cryptocurrencies however, means that one may need to gather data from multiple exchanges simultaneously, matching transactions to previous breeding events where possible, or guessing the approximate values for payouts which do not appear on the gathered exchanges. This step is considered out of scope for this study given the focus on the exclusive use blockchain transactions, but presents an interesting and challenging area of future work.

Of the set of measures which fall under the time and cost domains, only those that are also in the subset of measures used in early work [61][2] are applied (again given the youth and sparsity of behavioural measures used as described in Chapter 5), as they typically concern high level behavioural descriptors such as the total amount wagered, and total number of bets, which aim to capture levels of temporal and financial gambling involvement. This means they can be more broadly compared to existing studies in the player behaviour tracking domain, as opposed to more recently created

measures which may not be broadly comparable. The set of behavioural measures applied in this study are duration, frequency (percent), number of bets, mean bets per day, mean bet size and total amount wagered. These were chosen because they have each been broadly applied in the field of player behaviour tracking and can be (tentatively) related in different experimental configurations to external harm related variables such as self reported harms due to problem gambling (See Chapter 5).

Each of these measures will be computed across a collection of transactions from the CryptoKitties `breedWithAuto` function call transaction set stored on the Ethereum blockchain. Further analysis will explore Spearman's correlation coefficients across each of the pairs of measures in this set. This two-step approach follows a methodology common to studies in the field of player tracking [2][94][105], and will contribute to our understanding of the currently unknown relationships between these measures in this technologically advanced, and gambling-adjacent domain. Importantly, this will allow not only a comparison of each measure individually at the population level, but of how the measures relate to one another in comparison to gambling activities.

The third distinct analysis will be the isolation of so-called 'heavily involved players', which can be extracted as the top percentage of players by total amount wagered [2]. As discussed in the systematic review of behavioural measures used in player tracking research (See Chapter 5), this division along a single measure is not meaningful in distinguishing between at-risk problem gamblers, but rather acts as an epidemiological descriptive technique for understanding the distribution of levels of involvement within the population. This is important because not all players who exhibit heavy involvement are at risk of problem gambling, but a large portion of players who exhibit heavy involvement have been linked to self identified gambling related problems [3].

Data Sample

The sample of successful transactions to the `breedWithAuto` function call in the CryptoKitties application consisted of 1,629,171 transactions originating from 64,044 unique addresses. This sample's first successful transaction occurred at 18:57 on 23/11/2017, with the last at 04:48 on 06/08/2021, meaning the total sample covers three years and eight months, or a full

1352 days. Like previous studies in player behaviour tracking research, this study clips players whose durations (time between first and last interaction with the application) are less than three days. This acts to remove players who have simply tried out the application and then quit, thereby skewing data. Of the 64,044 unique addresses (players) in the sample, 21,897 (34%) engaged with CryptoKitties for three or more days, so were taken forward to the following analysis.

8.3.3 Results

Population Behaviours

The behavioural measures computed across the entire (n=21,897) sample are presented in Table 8.1. In order to make these results comparable with the centralised and decentralised gambling literature, parametric and non-parametric measures of centrality are presented, plus results of a one sample Kolmogorov-Smirnov test which tests for normality (the p-values for which are in column p).

The results of the K-S tests indicate a highly non-normal distribution across all measures. This is to be expected across behavioural measures, and is a phenomena well-documented in existing literature [1]. Confirming this in the blockchain games domain means the parametric statistics presented can be largely ignored, however they remain included to provide a reference for the non-parametric equivalents. For example, it is clear that the number of bets is likely non-normal, and with median and inter-quartile range value of 7 and 12 respectively. However, the mean and standard deviation values of 69 and 1734 show that the sample indeed exhibits massive skew, the extent of which is not clear from the non-parametric values alone.

The first measure of duration, given the more than three year observation period, appears to be very low, with a median player interaction period of 11 days (See Figure 8.3 for centile distribution). Furthermore, even with the upper inter-quartile range considered, an engagement duration of 47 days out of a possible 1352 shows that the majority of players generally are not particularly invested in this application. This observation is matched in the number of bets, which with values of 7 and 12 as above, shows that the majority of players do not engage very heavily in terms of their interactions with this particular (`breedWithAuto`) mechanism.

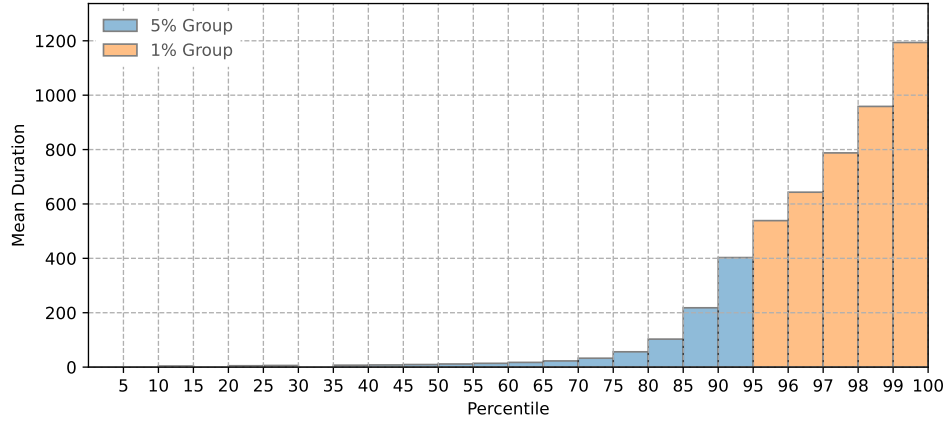


Figure 8.3: Distribution of the behavioural measure of duration within the blockchain games player sample ($n=21,897$). Note that the x-axis is discontinuous to highlight the extreme nature of the distribution.

Table 8.1: Behavioural measures applied to CryptoKitties' `breedWithAuto` function calls.

	Mean	Std	Median	IQR	K-S Score	<i>p-value</i>
duration (days)	88.53	201.79	11.00	36.00	1.00	< 0.01
frequency_percent	39.81	30.21	36.43	47.23	0.89	< 0.01
number_of_bets	68.97	1734.79	7.00	12.00	0.98	< 0.01
mean_bets_per_day	3.15	7.14	1.73	1.83	0.84	< 0.01
mean_bet_size (ETH)	0.01	0.01	0.01	0.0	0.50	< 0.01
total_wagered (ETH)	0.58	13.96	0.06	0.12	0.50	< 0.01

Table 8.2: Spearman’s correlation coefficient scores between each pair of the six measures computed across CryptoKitties `breedWithAuto` function calls. Key: `dur` = duration, `freq` = frequency percent, `num` = number of bets, `mday` = mean number of bets per day, `msize` = mean bet size, `tot` = total amount wagered. All values are significant at $p < 0.05$.

	dur	freq	num	mday	msize	tot
dur	-					
freq	-0.81	-				
num	0.24	0.23	-			
mday	0.05	0.26	0.87	-		
msize	-0.19	0.19	-0.04	-0.03	-	
tot	0.2	0.24	0.93	0.81	0.24	-

Understanding the typical temporal involvement of players is then complemented by the use of frequency percent and mean bets per day, which both show that although players may engage a median 36% of days, placing a median 1.73 bets. This pattern varies dramatically from the decentralised gambling domain [87], which (even including players whose durations are less than 3) place a median 6 bets per day. Finally the financial involvement oriented measures of mean bet size and total amount wagered again diverge from what we now know in the decentralised gambling domain. In the case of mean bet size this is to be expected, as the ‘bet’ in this context is fixed by the application’s developers so therefore is not expected to change. The total amount wagered however is dramatically lower, with a median and inter-quartile range of 0.06ETH and 0.12ETH respectively. This is extremely low in comparison to the gambling domain (See Chapter 7), and indicates a generally low level of financial involvement amongst players with this mechanism.

As in existing studies of gambling behaviours, exploring the rank order correlations between each of the different behavioural measures provides insight into the relationships between gambling behaviours [100]. Table 8.2 shows the Spearman correlation coefficients and their significance between each pair of behavioural measures computed. The largest coefficient of 0.93 between total wagered and total number of bets makes sense in the context of the `breedWithAuto` function as the ‘bet size’ or value is fixed by the application itself. The reason this correlation is not perfect can be explained by changes to the fixed value by the developers over the lifetime of the smart contract. Similarly, strong positive correlations between mean bets per day

Table 8.3: Descriptive statistics on the set of behavioural measures computed (n=21,897), split by heavy involvement (top 5%) by total amount wagered. Key: dur = duration, freq = frequency percent, num = number of bets, mday = mean number of bets per day, msize = mean bet size, tot = total amount wagered.

	Top 5% (n=1,095)				Other 95% (n=20,802)			
	Mean	Std	Median	IQR	Mean	Std	Median	IQR
dur	290.37	350.89	133.00	445.50	77.91	184.75	10.00	28.00
freq	47.42	34.18	37.80	65.39	39.41	29.94	36.36	47.50
num	1132.28	7683.42	255.00	491.00	13.00	20.64	6.00	11.00
mday	17.62	26.11	11.07	13.40	2.38	2.50	1.67	1.67
msize	0.01	0.00	0.01	0.00	0.01	0.01	0.01	0.00
tot	9.40	61.80	2.25	4.10	0.12	0.15	0.06	0.10

and total number of bets, and between total wagered and mean bets per day, make intuitive sense, as they capture aspects of the volume and rate of financial involvement with this particular mechanism in the application.

Aside from these obvious strongly positive correlations, the negative 0.81 between frequency percent and duration echoes engagement with gambling applications, but as with all applications, the greater a player's duration the less likely they are to engage with it every day generally, so this is of less phenomenological interest. All other correlations are no greater than 0.26 so do not appear to be meaningful in any descriptive sense so are not discussed further.

Heavily Involved Players

Taking the entire population's behavioural measures described above and partitioning them by membership in the top 5% by the measure of total wagered yields the split descriptive Table 8.2. This table tells a very different story of the population, with this top 5% group exhibiting dramatically higher values across all behavioural measures except frequency. As in the previous Study in this thesis, Table 8.4 shows LaBrie et al's original results [2] for quick comparison.

This heavily involved group (n=1,095) engaged for a median 133 days, placing a median 255 'bets' at a rate of 11 per day. This higher rate of spending, paired with the fixed 'bet size' is reflected in a median and inter-quartile range in total wagered of 2.25ETH and 4.10ETH respectively. While these values are still not close to their gambling counterparts, they

Table 8.4: Gambling behaviour of extreme 5 and 95% subgroups of casino bettors, reprinted from LaBrie et al's 2008 study [2].

	Most involved casino bettors top 5% ($n = 212$)		Other 95% of participants ($n = 4,010$)	
	Mean (SD)	Median	Mean (SD)	Median
Duration	476 (232)	529	290 (233)	246
Frequency	24% (17)	20%	16% (21)	7%
Number of bets	24,558 (36,779)	10,465	2,403 (7,819)	486
Bets per day	285 (344)	188	107 (176)	46
Euros per Bet	213 (682)	25	25 (97)	4
Total Wagered	345,579 (354,890)	233,195	10,338 (19,360)	2,284
Net Loss	8,746 (11,213)	6,698	422 (939)	107
Percent Loss	2.6 (3)	2.5	8.0 (12)	5.9

do highlight an almost 40x increase compared to the remaining 95% of the population. It is however important to remember that this 5%-95% split is on the 34% of the total population whose duration is greater than 3, making the top 5% ($n=1,095$) here equate to the top 1.7% of the total population who have ever engaged with this mechanism.

8.3.4 Discussion

The results presented above provide a population level behavioural profiles of blockchain gamers who engage with the CryptoKitties breedWithAuto mechanic. Beginning with the data sample itself, this study found the size and longitude of the data available on the Ethereum blockchain available for this type of analysis to be extensive. This, even when a duration-based exclusion criteria was applied, left transactions from over 20,000 unique addresses open to analyses.

Using this large sample size, the behaviours at the population level presented a somewhat underwhelming picture of engagement with the randomised reward mechanism within the CryptoKitties application. This is in stark contrast to recent studies in games research which have found a link between problem gambling severity and the use of randomised reward mechanisms [127][128], which may intuitively suggest high engagement with such mechanisms. Although many of the behavioural measures computed were heavily skewed, the median and interquartile ranges across the population appear low, and not indicative of any heavy involvement or other patterns of phenomenological interest. The results of the Spearman correla-

tion coefficients between pairs of measures at the population level echoed this finding, with no unexpected correlations between measures of different types (time domain vs cost domain). These results, in the broader context of the known similarities between randomised reward mechanisms and forms of gambling, suggest that while there may be a relationship of some kind, it may not manifest in users actually engaging with such mechanisms in the same way as gamblers engaging with casino games.

The results of this study have also presented the behavioural profiles of a heavily involved subset of players, whose total amounts wagered are much greater than the broader population, but whose spending still appears to be much lower than in the decentralised gambling domain. The behaviour of this subset is indeed more extreme than the majority of players, but still generally less extreme than gamblers in the decentralised domain. This study is therefore cannot conclusively establish a behavioural link between randomised reward mechanisms and gambling in the cryptocurrency domain, although given the dangers that such an empirically supported [127][128] link poses in the centralised domain, further research in this area is required.

It is possible that the behavioural measures are different, particularly in the cost domain, due to the difference in the form of the payout, as unlike in casino games a player receiving a valuable reward from a randomised reward mechanism is not instantly able to reuse that payout for subsequent spending. As measures in the cost domain typically aggregate spending, the re-use of payouts can dramatically inflate measures such as total amount wagered and mean bet size. However, this need to liquidate the payout and resulting deflation of behavioural measures cannot explain the links repeatedly established in existing literature between randomised reward mechanism spending and problem gambling severity [127][128]. A second possibility then emerges that the profiles of blockchain gamers in particular show behaviours which are unique to their centralised counterparts. The lack of behavioural data in the centralised domain makes this possibility particularly difficult to explore, but as industry collaborations with academia grow, studies of centralised randomised reward mechanism transactions will no doubt appear for future comparison.

Limitations

The primary limitation of this study's methodology is that behavioural measures in the loss domain have not been applied. This is due to the difficulty in reliable valuation of the resulting NFTs at the moment of generation, paired with the rationale that this valuation may likewise not be immediately obvious to the player so may not be relevant. Interestingly, the low total amount wagered exhibited by blockchain games players in general, suggests that players are not selling the results of their randomised reward mechanism interactions in order to purchase more randomised reward mechanisms (thereby converting the output back into the original currency of the 'bet'), although this suggestion requires further empirical analysis to confirm.

A second limitation of this study is that it only considers transactions from the popular blockchain game CryptoKitties. This limitation is due to the time resources required to audit the architectures of blockchain games for specific mechanisms of interest, as to analyse every possible blockchain game would require several years. In CryptoKitties for example, the `breedWithAuto` function can be easily recognised and isolated, although this is not always the case. The use of a single blockchain game in this study means that the results may not generalise fully across blockchain games. This does not mean these findings are not informative, but rather that this study represents a first step into understanding behaviours in the emerging domain of blockchain games, and that further study is needed before any concrete conclusions can be drawn.

A third limitation of this study is that the data covers transactions over a period of several years. While this may not appear to be a limitation, in the context of the rapidly changing landscape of cryptocurrency adoption and the development of new games, behavioural data even a few years old may be considered out of date. This limitation can be overcome by the development of more sophisticated real-time analytics applications, although these are considered out of scope of this work.

Finally, in the decentralised gambling domain, differences in behavioural profiles between games have been used in order to identify and remove likely non-human engagement with given applications. In the case of engagement with the `breedWithAuto` function presented here, no such profiles for other games exist, inviting the possibility that no non-human players exist in this

data set, or that all of the players are non-human. While the former is unlikely, it is possible that a number of non-human accounts are skewing the data, although the unlike in the decentralised gambling domain, it is unclear what incentive may prompt the creation of such accounts. Nevertheless, exploring non-human engagement in blockchain games represents an interesting and related area of future work which may help address this limitation.

8.3.5 Conclusion

This study has presented an analysis of player behaviour when using a randomised reward mechanism in the popular blockchain game CryptoKitties. The methodology for this analysis has been informed by the field of player behaviour tracking - a branch of gambling studies concerned with understanding player behaviours. The analysis found that while a few high level behavioural similarities exist between player behaviour in blockchain games and decentralised gambling applications, these similarities appear superficial and importantly diverge when the scales of the total amount wagered are considered. In gambling research, both temporal and financial involvement have been found to be important aspects of problematic engagement. The generally low scores across behavioural measures used to capture both of these types of involvement with the `breedWithAuto` randomised reward mechanism show that while these mechanisms do share structural similarities with casino games, these similarities do not appear to be reflected in how users actually engage with them. This finding is heavily nuanced and may only apply to the blockchain games domain, but this study has applied one methodology by which such comparisons can be made.

Future Work

This study presents several areas of future work in both the blockchain games and decentralised gambling domains. Most obviously, the inconclusive similarities between player interactions with the randomised reward mechanism selected here and decentralised casino games warrant further investigation. This investigation could employ a different set of behavioural measures entirely, or focus on a particular domain of behavioural measures to pick apart further similarities and differences. Of particular interest are the trajectory based measures described in Chapter 5, which can be used to

uncover an escalation in involvement (or not) by players.

Conversely, existing work which associates problem gambling severity with randomised reward mechanism spending may be revisited with a view to obtain actual transaction data from affected players. The methodology applied here could then be replicated and the differences between the use of these mechanisms in blockchain games and centralised games could be better understood. Such pairing of qualitative and quantitative data could provide a holistic understanding of behaviours, leading researchers to better understand this relationship in the cryptocurrency domain and others. Whilst this study has found a lack of similarities between decentralised gambling applications and blockchain games spending behaviours, whether this difference generalises to their centralised counterparts remains unknown.

8.4 Summary

This chapter has temporarily departed from the central theme of gambling presented throughout this thesis. This departure is warranted by architectural similarities between specific mechanisms within blockchain games and gambling mechanisms found in decentralised gambling applications. An exploratory study through the lens of player behaviour tracking has revealed a very different landscape in player behaviours in the blockchain games domain than in the decentralised gambling domain. Their differences highlight a need for further research in the blockchain games domain, with a specific focus on their similarities to other forms of randomised reward mechanisms. Additionally, the transaction data centric approach used here can by definition only capture the actions of the players themselves, rather than the circumstances and environments they exist within, or their motivations for doing so. This limitation is inherent to all player behaviour tracking research, and does not prohibit analyses of this nature, but rather incurs a strict upper-limit to the generalisability of these results. As presented in the thesis overview in Chapter 1, this chapter concludes the exploratory branch into blockchain games.

Chapter 9

Identifying Behavioural Groups

“Science is what we understand well enough to explain to a computer. Art is everything else we do.”

Donald Knuth

Chapter 7 provided the first population level analysis of decentralised gambling application players using a widely used group of behavioural measures discovered in Chapter 5. This was a strong first step towards understanding behaviours in this emerging domain, however Chapter 5 also found that the analysis of any single behavioural measure in isolation has limited effectiveness when describing a population with respect to the potential gambling related harm they may experience. This limitation is overcome by the application of analytical methods which consider multiple behavioural measures simultaneously. One such approach used in existing work is player behaviour clustering described in Section 5.3. This chapter builds on the population level analysis of gambling in the cryptocurrency domain presented in Chapter 7 by applying player behaviour clustering techniques and comparing the unique behavioural groups uncovered in this domain to existing studies in the field.

This chapter presents the final study in this thesis. Section 9.1 presents the study itself, followed by a summary Section 9.2 which reviews the key findings ahead of the final chapter of this thesis. The broader implications of

the findings of this final study are deferred to Chapter 10, which will include discussion of all studies and findings uncovered by the research presented in this thesis.

9.1 Study 6: Behavioural Groups in Decentralised Gambling

In order to more accurately assess the scale of *potential* gambling related harm in the data sample gathered for this thesis, a more detailed exploration of player behaviours is required. Specifically, the nature and prevalence of different behavioural groups which have been empirically associated with a risk of experiencing gambling related harms can be explored, allowing direct comparison with existing studies in other forms of gambling. While this comparative approach cannot conclusively determine the number of players who actually experience gambling related harms, it can provide an estimate of the number of players *at risk* of harm when framed in the context of existing studies. This is important both for regulatory discussions around this emerging technology, and for informing the direction of further academic research. This study is therefore dedicated to answering all parts of Research Question 5 from the first chapter in this thesis: Which behavioural groups exist in the cryptocurrency gambling domain?

9.1.1 Introduction

The goal of this study is to identify unique behavioural patterns displayed during the first month of gambling in order to compare these patterns with other populations of gamblers. This selection of the first month of play is used as a cutoff to make between-player comparisons meaningful [3], rather than using a player's entire career data which may be of different lengths. A growing body of literature has explored the link between different behavioural measures and account closing due to gambling related problems at different periods in their careers [92][3]. These studies each take a labelled data set of players, with each label corresponding to their status as having closed their account, the reason for closure, plus a collection of transaction data containing their bets and payouts over a given period. From this data, a set of behavioural measures is computed, to which an unsupervised clustering

algorithm is applied in order to classify the population into distinct and behaviourally unique groups, under the assumption that those who experience different forms of gambling related harm engage with gambling applications in distinct ways. These clusters, and their membership, can then be compared to the labels such that relationships between certain behavioural profiles can be, or not be, associated with certain labels. For example, in Braverman and Shaffer's 2010 study of high-risk internet gambling [3], a group of players exhibiting high activity and high bet size variability was found to contain 73% players who reported closing their accounts due to gambling related problems. Similarly, Xuan and Shaffer [92] found that players who closed their accounts due to gambling related problems showed an increase in monetary loss, increase in stakes per bet, *and* increasingly shorter odds up to the time of their account closure.

This study aims to anchor its approach against Braverman and Shaffer's 2010 study on the identification of high-risk internet gamblers using their transaction data [3]. This study was chosen over others uncovered by the systematic review presented in Chapters 4 and 5 for four key reasons. The first is that it focuses on an open data set of live action sports bettors - a subject which has received a lot of attention in the studies uncovered by the systematic review presented in Chapter 5. This means that a much broader description of the players in this data set exists than any other comparable study, making it an excellent candidate for comparison. The second is that although this study is now over a decade old, the original author's colleagues at the Division of Addiction and Harvard Medical School have recently (2021) published a further study assessing the changes in this domain over the last decade [125]. Importantly, they find that 'sports wagering behaviour has remained relatively stable over time despite legislative changes and an increase in popularity'. This means that even though the original study's findings are now relatively old, they are still of relevance in light of the stability of behaviours in the sports wagering domain. The third is live action sports betting and casino gambling are mechanically relatively similar activities (depending on the time scale of the sports betting). The exact differences in terms of the behavioural measures between the two is still open to debate, with Ukhov et al's recent comparative work [53] being an insightful albeit tentative comparison (using a model with precision of 0.45 and 0.60, and recall of 0.27 and 0.42, for casino gamblers and sports bettors

respectively¹). A fourth reason that this study was chosen, is that no similar unsupervised machine learning studies exist on casino game transaction data.

K-Means Clustering

Braverman and Shaffer applied the k-means clustering algorithm [131] to a set of behavioural measures derived from player's first month of gambling activity. This algorithm is widely recognised as a simple and easily interpretable clustering method which seeks to partition a collection of observations into a predefined number (k) of clusters [132]. It does this by randomly adding k points to the data set, and iterating over their positions until their distances from other nearby data points becomes stable. These k points are known as centroids, and their positions are used as descriptors for the group of data points which they are closest to. This iterative partitioning relies on the knowledge of the value of k, which can be empirically determined by repeatedly performing clustering over a range of values of k whilst measuring some goodness of fit metric. As Kodinariya and Makwana describe [133], many goodness of fit metrics exist, one of the simplest of which is plotting a cost function against a range of values of k. One such cost function is that of inertia, which is simply the sum of the squared distance between each cluster centroid and its members. Plotting the mean inertia across a range of values of k can therefore be used to indicate the point at which the addition of new clusters only marginally improves the distance between each cluster centroid and its members. The 'elbow' of this plot can then be used to visually select acceptable values of k.

Research Questions

The behavioural profiles of subgroups within the cryptocurrency gambling population are unknown, which means the prevalence of behavioural groups which have previously been linked to harm related variables is also unknown. The volatile nature of cryptocurrency prices suggests that those who own them may exhibit a higher risk tolerance than the general population [134]. One may therefore expect a higher portion of cryptocurrency gamblers to be in an empirically determined high activity, high variability subgroup than in

¹This means that it only has a moderately reliable detection mechanism (0.45), and correctly detects only 27% (0.27) of true positives for casino gamblers.

other forms of gambling - a subgroup identified as containing mostly (77%) at-risk players by Braverman and Shaffer [3]. This study therefore explores the final research questions presented in Chapter 1;

- Which behavioural groups exist in this domain?
 1. How prevalent are they within the population?
 2. How do they compare to other forms of gambling?

9.1.2 Method

Data Sample

The data sample used in this study is an updated version of the data gathered as part of the fourth study in this thesis (Chapter 7), including more recent transactions in order to maximise the sample size of players. This includes transactions to and from the Dice2Win, Etheroll, and FCK decentralised gambling applications, which each correspond to the placement of a bet (each application's architectures are described in Section 7.2 above). As discovered in Chapter 7, the raw transaction data available on the Ethereum blockchain to and from these applications contains a number of likely non-human actions. These are removed from the data set by identifying each of the addresses which have at least one transaction to the FCK coin flip game, and removing all of their transactions. This removes a total of 598,327 transactions from a total sample of 2,934,795 leaving 2,336,468 for analysis. These approximately 2.3M transactions originate from 14,462 players, and cover a time period of almost three years; from August 2018 through to July 2021. The cumulative value across each of the respective applications over time is provided in Figure 9.1.

Measures

The measures computed for this study fall into two groups. The first group is the four variables which describe the pattern of gambling activity during the first month of play. These, as in the Braverman and Shaffer study [3], include the frequency (total number of active days), intensity (total number of bets divided by frequency, or mean bets per active day), variability (standard deviation of the total amount wagered per active day), and trajectory (the trajectory of total amount wagered per active day). The original author's

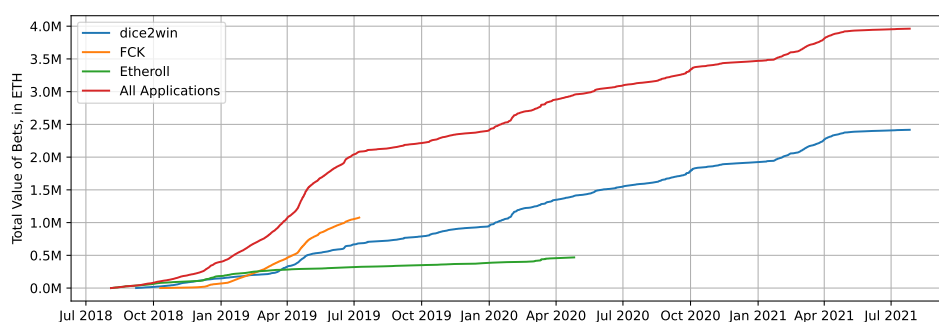


Figure 9.1: Total value transacted (ETH) by all bets placed in the data sample ($n=14,462$, $tx=2,336,468$).

Table 9.1: Behavioural measures which are used as input to the clustering algorithm (as used by Braverman and Shaffer [3]), computed using players' first month of transaction data.

Name	Description
Frequency	Total number of active days
Intensity	Total number of bets divided by frequency
Variability	Standard deviation of total amount wagered per active day
Trajectory	Coefficient of total amount wagered per active day

rationale for selection of these four measures in particular is based in findings from several previous studies which relate each of them to uncommon behaviours within a population [2][88][61][100], plus an operationalisation of the ‘need to gamble with increasing amounts of money’ described in the American Psychiatric Association’s Diagnostic and Statistical manual of Mental Disorders [135] for gambling disorder. Table 9.1 provides a summary of these measures for reference. These measures will each be standardised using the z-score transformation as in the original study. This transformation will enable direct comparison of cluster centroids with the original study, as z-scores describe an observation’s position relative to the population mean.

The second group of measures are those which describe the players over the entire duration of their careers to date. Again as in the original study, this includes the period of gambling (duration), total amount wagered, total number of bets, mean bets per day, and the net loss for the entire period. These variables were originally chosen to ‘summarise betting behaviour for the entire period of gambling’ [3], and act in the same way as in Studies IV and V in this thesis to provide richer descriptions of player behaviours. An

Table 9.2: Behavioural measures which are used to describe player’s behaviours, computed on both the first month of data and entire careers of data - these are not used as input to the clustering algorithm.

Name	Description
Total Wagered	Sum of all bet sizes
Number of Bets	Total number of bets placed
Mean Bet Size	Total wagered divided by number of bets
Duration	Calendar days between first and last bet
Net Loss	Total wagered minus total payouts

additional computation of this second group of measures across transactions from each individual’s first month of play will also be performed. This was not performed in the original study, however, because this study uses unlabelled data, this first month description may help better contextualise the behaviours shown in each of the clusters, providing greater context for the results of the clustering. This dual month and career descriptive approach also allows comparison between each player’s first month (the transactions on which they were clustered), and their career. This comparison may reveal important details about the way in which different clusters engaged with these gambling applications over time.

Clustering

In order to make the results of this study directly comparable to those of Braverman and Shaffer’s existing work [3], the k-means clustering algorithm will be applied, as this was applied in the original work. Following correspondence with the original authors, the original analytical code is not available, so a readily available implementation of the algorithm from the scikit-learn library will be used [136]. Although the original study finds a four cluster solution to be most ‘stable and reliable’, I do not assume that this solution holds in the cryptocurrency domain. For this reason, an ensemble of solutions will be computed, with the mean inertia score across each k value plotted against k. Here, the inertia score is the total absolute distance between each cluster centroid and its member data points, which naturally decreases with increasing k. This plot should exhibit an ‘elbow’ much like a scree plot in principle component analysis, which occurs as the rate of decreasing inertia slows as k increases due to new solutions more completely capturing the ‘true’ clustering in the data set. The elbow in this plot can be used to inform

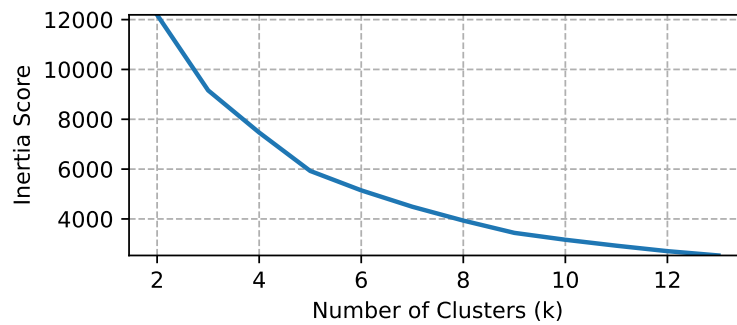


Figure 9.2: Mean inertia scores across repeated clustering solutions in the data sample (first month of gambling, repeats per $k=30$, $n=3,870$).

the final value of k , as an extreme solution would have an inertia of zero and a value of k corresponding to however many distinct groups exist within a given population (under the assumption that groups are unique and contain homogeneous members).

9.1.3 Results

Cluster Solutions

Figure 9.2 presents the mean inertia scores across the solution ensemble for different values of k . Unlike Braverman and Shaffer's original study, this figure shows that a five or six cluster solution appears to be the point at which the rate of inertia decrease with increasing k begins to slow. This means that although the inertia score generally decreases with increasing k , at values of 5 and 6 for k the rate of decrease is still higher than this expected decrease. Although there is rarely a single 'correct' answer as to what the value of k should be for any given data set, the value of 5 is taken forward for the remainder of this study. This limitation is discussed further in Section 9.1.4.

This choice of 5 for the value of k represents a margin call driven by two forces. On one hand, the value of k should be as low as reasonably possible in order to segment the sample into approximate groups, not an 'over-fit' solution which may classify one part of a larger group as its own distinct group. On the other, the value of k should be as high as reasonably possible to ensure that any groups in the population are successfully identified, and that no similar but distinct groups are incorrectly clustered together. It is

Table 9.3: Standardised behavioural measure (number of standard deviations from mean) scores of cluster centroids using a five cluster solution across the first month of cryptocurrency gambling activity (n=3,870). Extreme values are displayed in bold.

Cluster Number	1	2	3	4	5
Cluster Size	n=3092	n=458	n=279	n=29	n=12
Description	moderate betting	high frequency	high intensity	extreme variability	negative trajectory
Frequency	-0.35	2.16	0.32	0.50	-0.51
Mean Bets per Day	-0.27	0.24	2.61	-0.35	0.68
Bet Size Deviation	-0.09	-0.06	-0.01	10.08	2.44
Bet Size Trajectory	0.02	0.04	0.18	0.43	-12.10

Table 9.4: Clustering results reproduced from the original study for reference [3]. Extreme values are displayed in bold.

Cluster Number	1	2	3	4
Cluster Size	n=15	n=22	n=115	n=378
Description	high activity, high variability	low first-month activity	high activity, low variability	moderate betting
Frequency	2.63	-0.54	2.39	0.28
Mean Bets per Day	1.79	0.04	1.90	0.00
Bet Size Deviation	4.41	0.16	0.26	-0.04
Bet Size Trajectory	0.27	-2.49	0.14	0.22

important to note that this is also in the context of providing an approximate description of groups within the population, rather than a robust classification of some pre-existing label as typically done in studies of machine learning. I can therefore proceed with a value of 5 in this study, although as with any descriptive technique, further studies may pursue higher or lower values.

Descriptions of Clusters

Table 9.3 shows the behavioural profiles of each of the groups in the five cluster solution, plus the number of members of each. As in Braverman and Shaffer’s original study (results reproduced in Table 9.4 for reference), a single majority group exists which exhibits near average scores across each of the four behavioural measures. This is proportionally larger (80%) in comparison to the 73% majority group in the original study, leaving 20% or 778 players across the remaining four clusters combined. Clusters 2 and 3 are the second and third largest by member count respectively, and each exhibit

Table 9.5: Behavioural measures describing the first month of gambling activity for each cluster identified using the k-means algorithm. Notable values are displayed in bold.

Cluster Number	1	2	3	4	5
Cluster Size	n=3,092	n=458	n=279	n=29	n=12
Total Wagered	82.52	609.47	2,735.44	14,266.92	3,109.19
Number of Bets	96.53	873.15	1,462.65	149.07	329.58
Mean Bet Size	0.82	0.76	1.12	112.56	27.45
Duration	13.72	26.36	14.87	16.72	9.33
Net Loss	1.22	-12.04	15.30	220.33	61.07

extreme behaviour along a single behavioural measure (in bold in Table 9.3). Cluster 2 contains 458 players whose frequency is much higher than any other cluster, but whose other measures remain close to the population average. Cluster 3 contains 279 players, although their mean bets per day (intensity in original study) was higher than any other cluster - again with other measures close to the population average.

Clusters 2, 3, and the majority group, account for 3,829 (99%) of players in the data sample, leaving 41 players between clusters 4 and 5. These clusters each exhibit extremely high scores along a single behavioural measure, making them significant outliers from all other groups. Cluster 4, with 29 members, exhibits a bet size deviation score of 10.08 yet somewhat counter-intuitively exhibits typical values across the remaining three measures. Cluster 5, with just 12 members, shows a bet size trajectory score of -12.10, plus a high bet size deviation score of 2.44. This highly negative bet size trajectory means that players in this group are rapidly decreasing the sizes of their bets, which also partially explains the high bet size deviation score. Using these cluster descriptions alone, it is clear that three large groups have been identified by the k-means algorithm, with a further two capturing the few players who do not fall into one of these three groups - this mix of three large and two small will naturally be affected by the choice of k .

Month Descriptions

Table 9.5 shows the mean behavioural measures for each of the cluster centroids across the first month of play. These (month descriptions) were not computed in the original study, but provide richer insight into the behavioural differences between each of the clusters which Braverman and

Shaffer's approach yield. The first obvious difference between each of the clusters is that although none of the four measures on which they were clustered contain any information regarding the sum of wagers, the total amount wagered varies dramatically between the five clusters - with the exception of clusters 3 and 5. For example, in the first month of activity, players in the majority group typically wagered a total of 82.52 ETH, whereas those in cluster 2 wagered a total of over 600 ETH, an almost 8x increase which continues through cluster 3 (2,735 ETH) and 4 (14,267 ETH). Cluster 5's total amount wagered of 3,109 is closest to that of cluster 2, although it deviates substantially in mean bet size and duration, standing out as a group of 12 players who placed many bets, of substantial size (27.45 ETH), and who lost the second highest amount of any of the clusters (61 ETH).

This first month descriptions adds to the cluster descriptions in Section 9.1.3 by highlighting that the 279 players in cluster 3, despite showing only abnormal intensity (mean bets per day) along the clustered measures, show both high total amount wagered and a high number of bets in comparison to the other clusters. This heavy financial involvement paired with the high temporal involvement that placing a mean 1,462 bets over a period of 15 days incurs, indicates that although this subgroup does not match Braverman and Shaffer's high activity, high variability subgroup (and therefore cannot conclusively be related to any harm related variable), it may still warrant further investigation. Similarly, clusters 4 and 5 when viewed through the clustered measures alone do not show high activity and high variability, but as with cluster 3 they show a large number of bets, large net losses, and high total amounts wagered, making them candidates for further work in understanding the external conditions of their members.

Career Descriptions

Table 9.6 presents the behavioural measures of members in each of the clusters, computed across their entire gambling careers with the three applications. As with the first month descriptions in Table 9.5, this information can be used to create a more detailed picture of the different behavioural profiles of each of the clusters, yet over the course of each player's entire career. For example, the career level behavioural measures in tandem with the first month measures above show that players in cluster 5 typically only play for 12 days, therefore showing similar characteristics at both the one month

Table 9.6: Behavioural measures describing the entire careers of gambling activity for players in each cluster identified by the k-means algorithm. Notable values are displayed in bold.

Cluster Number	1	2	3	4	5
Cluster Size	n=3092	n=458	n=279	n=29	n=12
Total Wagered	151.92	1029.57	4563.19	23073.17	3111.24
Number of Bets	181.75	1595.82	2075.66	650.97	341.92
Mean Bet Size	0.83	0.75	1.16	116.26	27.43
Duration	75.16	148.65	87.02	60.00	12.25
Net Loss	1.95	-4.89	39.68	274.39	61.10

and career scales. The duration scores for players in cluster 2 echo the first month’s activity, as with a mean duration of 148 days they are the group engaged with the three applications over the longest period, and again somewhat counter-intuitively are the only group whose net loss is negative - although not as negative as in their first month of play.

Further behavioural differences at the career level exist between players in clusters 2 and 3, as although they place a similar number of bets (1,596 versus 2,075), those in cluster 3 wager over 4x the amount, and typically lose 40 ETH in comparison to cluster 2’s small profit. Those in cluster 3 also engage with the three applications for approximately half as long as those in cluster 2, making cluster 3’s high losses in a relatively short period a concerning and potentially fruitful area of future work. This view of cluster 3’s behavioural profile at the career level supports the concerning features presented at the first-month level.

9.1.4 Discussion

This study found that a five cluster solution using the k-means algorithm could capture behavioural differences between players in the cryptocurrency gambling domain. This differs from Braverman and Shaffer’s original study [3] which finds a four cluster solution of players to be stable and reliable in the live-action sports betting domain. This difference in the number of clusters between these two studies may be due to sampling differences, as this study used all available bets for each player, whereas the original study used a subset of players who had closed their account between 1 month and 2 years of starting playing. As the concept of a player ‘account’ in the applications analysed in this study does not apply, this eventual desire for exclusion from

gambling cannot be inferred without using a survey or similar data gathering method to augment the available transaction data set - this is expanded upon in the limitations section below.

The five clusters identified each exhibit different behaviours, although no single cluster matches the high activity and high variability behavioural profile identified in the original study. One can therefore reject both hypotheses posed in Section 9.1.1, that behavioural subgroups in the cryptocurrency domain match those found by Braverman and Shaffer, and that this subgroup is proportionally larger in the cryptocurrency domain. This is important as 77% of this group's membership in the original study was found to have closed their account due to gambling related problems. The lack of a similar distinct behavioural group in this data sample, despite Study 4 (Chapter 7) in this thesis finding a heavily involved subgroup spending more than in comparable populations of gamblers, means the scale of those who may be at risk of experiencing gambling related problems in the cryptocurrency domain remains unknown. This considered, one cluster identified in this study contains 279 (7% of players) members whose temporal and financial involvement across the applications presented is dramatically higher than the empirically determined 'majority group' (3,092 or 80% of players) in their first month of play. This higher involvement however was only detected by examination of descriptive measures computed across their first month of play - a technique not applied in the original study. While this result is therefore inconclusive with respect to understanding gambling related problems in the cryptocurrency domain, it has still been informative, and has highlighted significant deviations between distinct subgroups within the population. Understanding more holistically what causes these deviations, and possible relationships to other harm related variables remains an important area of future work.

One counter-intuitive result uncovered in this study is that a behavioural group exists in the cryptocurrency gambling domain whose typical member profits from their interaction with the gambling applications (despite no loss related variables being used in the clustering). This may simply be by chance, or may be related to how they engage with the applications; This groups duration of play was the highest of all groups, and their number of bets placed second highest. While this may simply be an artefact of any population of players engaging with casino games, none of the groups

identified in the original study exhibited such negative losses at the career level. This may imply that behavioural groups in casino games generally are distinct from live action sports bettors, or that behavioural groups in cryptocurrency casino games are distinct from live action sports bettors. This invites similar analysis of casino game players so that a deeper understanding of the behavioural differences between different types of gambling can be generated, which may support or refute Ukhov et al's [53] early analysis.

Limitations

The primary limitation of this study is that it only applies a single clustering algorithm - a limitation which at first seems trivial, but in the context of existing player behaviour tracking becomes somewhat complex. This is because although applying other clustering algorithms is an obvious and technically simple next step, no comparable studies exist in other (labelled or unlabelled) domains with which to compare results. As k-means was used in the original study by Braverman and Shaffer, it was also used here in order to make results directly comparable. This means that although several unique behavioural clusters were identified, it is difficult to know conclusively whether they represent the ground truth of behavioural partitions within the population, or whether they are simply one possible solution in a multiverse of possible solutions. This limitation could be overcome by applying multiple clustering algorithms, with a goal to explore how different behavioural subgroups within the population are identified by these different algorithms, and how stable these groups are across the different algorithms. This multiple algorithm approach could support the findings presented in this study, or be used to assess the reliability of clustering itself as an approach to understanding player behaviours in the cryptocurrency domain. However, as with any study in an emerging domain, such work would need grounding against existing studies or replication on existing data sets in order to generate meaningful insights. This represents an important but large area of future work.

A second limitation of this study shared by all studies in this thesis, is that the prevalence of those experiencing gambling related harms is unknown in the cryptocurrency transaction data sample. This means that while a comparable study was used to frame findings in an attempt to understanding its prevalence in the cryptocurrency domain, the rate of actual experience of

gambling related harms may be much higher or much lower in this sample than in the original study, thereby potentially skewing the results. For example, it could be that the majority of members in one or all of the clusters are experiencing gambling related harms, yet as no single cluster matches Braverman and Shaffer's high activity and high variability group, the estimated prevalence of these harms remains unknown. As with behavioural research in other domains, transaction data oriented studies alone cannot conclusively infer complex situations such as experiencing gambling related harms, or an understanding of a player's motivation to gamble. This does not mean that the finding presented here are not informative, but rather that they represent a first step towards building a holistic understanding of behaviours in this emerging domain. This highlights the need for both qualitative and quantitative data to be used together in further work to address this limitation.

A third limitation of this study is that only a select set of behavioural measures were computed based on the original study [3] which were then used as input to the k-means clustering. This means that although from this study it is clear that Braverman and Shaffer's method yields distinct behavioural groups, it is unclear whether or not their originally identified groups may actually exist in this sample but may be better identified in the cryptocurrency domain using a different set of behavioural measures. For example, the addition of a measure from the loss domain such as net loss may provide enough additional information such that a four cluster solution becomes preferable. Similarly, this study found that despite not using total amount wagered as measure for clustering, each of the clusters' values of total amount wagered were different, making it a potentially important explanatory variable for behavioural segmentation of players. Computing additional measures from any domain would however be difficult to justify, as it would be impossible to ascribe meaning to such results without findings from identical application in another domain. As with the first limitation above, this can be overcome by more studies in other domains using identical methodologies, from which a richer understanding of behavioural groups within the cryptocurrency gambling domain could be generated.

9.2 Summary

This study has applied the k-means clustering algorithm to actual cryptocurrency betting data derived from three casino game applications on the Ethereum blockchain, and presents the first ever analysis of behavioural groups within a population in this domain. The results of this clustering shows several differences to previous work in player tracking research, with a five cluster solution appearing optimal using inertia scores across an ensemble of possible solutions. This increased number of clusters in comparison to existing work revealed a single majority group containing 3,092 (80%) of the sample which contained no extreme behaviour, two moderately sized clusters of 458 (12%) and 279 (7%) each showing high frequency and high mean bets per day respectively, and two smaller clusters of 29 (<1%) and 12 (<1%) which showed extreme behaviours in bet size deviation and (negative) bet size trajectory respectively.

Descriptive behavioural measures across the first month and entire careers of play revealed extreme spending and losses in clusters 4 and 5, with cluster 3 showing heavy involvement by number of bets, high total amount wagered, and significant losses. The clustering methodology used in this study cannot conclusively relate membership with any of these clusters to an external harm related variable, but has identified these three groups' extreme behaviour, presenting a potentially important direction for future work.

This chapter has presented the final study in this thesis, which built upon the descriptive findings of gamblers in the cryptocurrency domain presented in Chapter 7 by applying a clustering approach used in existing work by Braverman and Shaffer [3]. This approach revealed five distinct behavioural groups, three of which presented extreme behavioural measures which in the context of findings of the systematic review presented in Chapter 5 present a potentially important area of future work. While this study's findings are inconclusive with respect to understanding the scale of potential at-risk players in this data sample, the results have still been informative, and provide a first step towards understanding behavioural groups in this domain.

Chapter 10

Conclusion

*“The world is changing, the planet’s heating up.
What the fuck is going on?”*

Bo Burnham
Comedy

This thesis has presented six studies, including a systematic review, two studies of the prevalence of decentralised applications containing gambling and gambling-like mechanisms, two studies of the behavioural distributions of users of a sample of these applications, and finally a study of behavioural groups within decentralised gambling applications. The key findings from these studies are;

- Only a handful of behavioural measures have seen enough use to reliably relate them to external harm related variables such as self reported gambling related harms (Figure 4.3).
- Several decentralised gambling and gaming applications offer large data sets for academic research (Chapter 6).
- A subset of players in the decentralised gambling domain show very high financial involvement (Chapter 7).
- Typical engagement with both decentralised gambling applications and randomised reward mechanisms in blockchain games is non-intensive (Chapter 8).
- Behavioural clustering of these gamblers was dissimilar to groups found in previous work, inviting further exploration (Chapter 9).

This chapter provides a brief summary of the research completed as part of this thesis, revisiting each of the research questions posed in the first chapter in light of the contributions made. It then discusses the implications of each of the key findings. Broader limitations of this research are then described, followed by a closing discussion of the future work which these findings suggest.

10.1 Contributions

10.1.1 Behavioural Measures

Study 1 in this thesis presented a systematic review of existing research in the field of player behaviour tracking, addressing research question 2; Which analytical methods can be meaningfully applied (to gambling transaction data). The key finding from this review was that the field is not yet developed enough for a robust theory of the relationship between specific behavioural measures and external harm related variables to have emerged. This significantly limits the application of behavioural measures to new domains, such as decentralised gambling applications, although as Studies 4 and 5 showed, this limitation is not prohibitive to understanding new domains.

This review also found that the behavioural measures used in existing work can be taxonomised according to the data required to compute them, and that using this taxonomy reveals areas (risk domain (Section 5.2.4) and miscellaneous domain (Section 5.2.5) which despite using intuitively important data (betting odds and house edge) have received little academic attention. While the work in this thesis did not set out to further develop these methods themselves, the finding that these gaps exist is important to further research in the field of player behaviour tracking, and is again a limiting factor in exploring new domains.

10.1.2 Decentralised Application Prevalence

Studies 2 and 3 in this thesis each explored the prevalence of decentralised applications, finding that in both gambling and gaming rankings, a small subset of applications are clear market leaders across several key metrics. These studies each contributed to answering research question 3 posed in the first chapter; How prevalent are decentralised gambling applications, and which

applications may provide the most data. The key contributions in each of these studies are the resulting selection of applications which present fruitful options for subsequent analysis, including *dice2win*, *etheroll*, *fck.com*, and *CryptoKitties*. Additionally, the finding that taking the classifications of any group of applications from a ranking service at face value will likely include several applications which are misclassified. In the case of blockchain games such misclassifications do not appear to adversely affect the market-level analyses, however several applications are incorrectly classified as decentralised gambling applications. These misclassified ‘gambling’ applications included large sets of transactions with a large number of ‘users’, thereby skewing analyses before their removal. The solution used to address this problem in this thesis was a manual exploration of the top portion of applications, however some form of automated smart contract classification presents an interesting and technically challenging area of future work.

10.1.3 Decentralised Gambling Behaviours

Study 4 applied the findings from studies 1 and 2 to provide the third important contribution in this thesis; a description of player behaviours in the decentralised gambling domain. This description applied behavioural measures broadly used across player behaviour tracking studies, finding that a small portion of heavily involved bettors in this emerging domain wager proportionally more than in other domains. This provides evidence to address research question 4; What are player behaviours in these applications, and how do they compare to other forms of gambling. Additionally, the typical user of decentralised gambling applications does not engage with the applications in any long term or financially significant way. This finding is somewhat counter-intuitive given the largely unregulated nature of these applications, and their lack of consumer protection mechanisms, but may be a feature of their simplistic and minimalistic design [137].

The application of behavioural measures to the transaction data generated by these applications also revealed a portion of players which exhibit suspiciously similar behaviours. Users of the *fck.com* application’s coinflip game in particular were significantly different across behavioural measures from the measures computed across other application-game combinations. This set of users can be tentatively labelled as ‘non-human’, as the behavioural profiles discovered could be created by a simple betting algorithm, although further

work on the deployment of automated betting algorithms to decentralised gambling applications is an interesting area of future work.

10.1.4 Blockchain Gaming Behaviours

Similarly to Study 4, Study 5 applied behavioural measures to transaction data gathered from the CryptoKitties blockchain game, specifically transactions to its `breedWithAuto` function, which acts as a gambling-like mechanism but instead of probabilistically returning currency, it probabilistically returns a kitten of unknown value. Unlike the decentralised gambling applications, this study revealed a pattern of much lower spending across the population of players, including a top portion of players whose involvement was much lower than in gambling applications. This generally lower involvement in blockchain games in comparison to decentralised gambling applications makes intuitive sense given the low payout liquidity, as the return of a kitten rather than currency means that players cannot instantaneously use payouts to place more ‘bets’. This finding does however invite further research to explore how the link between randomised reward mechanism engagement and problem gambling severity manifest in terms of actual behaviours in gaming applications.

10.1.5 Behavioural Groups in Decentralised Gambling Applications

The sixth and final study in this thesis applied a behavioural clustering approach established by Braverman and Shaffer [3] to the behavioural profiles generated using transaction data from the decentralised gambling applications uncovered in Study 2. This addresses research question 5; which behavioural groups exist in the decentralised gambling domain. Unlike in the original study, the data set taken from these applications is unlabelled, meaning that the rate of those at risk of gambling related harms is unknown. This study therefore applies a comparative approach, instead seeking groups whose behaviours are similar to those identified in the original study such that labels may be tentatively inferred. Unfortunately, the application of Braverman and Shaffer’s clustering approach did not yield any single group which exhibited similar behaviours to their ‘high activity, high variability’ group - 77% of which were at risk of gambling related harms. This inconclu-

sive result may be due to sampling differences between this study and the original, subtle differences at the implementation level of the approach as the original code was not available, or may be the ground truth of behaviours in the decentralised gambling domain. As with inconclusive results in other disciplines, this study invites more work into the application of other clustering approaches, surveys seeking labels for players in the data set, and more in order to confirm or refute this finding - however, such further studies will likely be very resource intensive.

10.2 Limitations

Each of the studies in connection with this thesis are subject to a number of limitations. In the case of Study 1, it is possible that the choice of search tool and databases may not fully capture all of the player behaviour tracking research which has been published. While the comparison against the results returned by my review with the studies found by Chagas and Gomes' review [1] supports the idea that the majority of studies have been identified, coverage of existing literature can always be improved by using more databases, more languages and so on.

Studies 2 and 3's limitations can be broadly grouped together as they share the same methodology. The first limitation of this methodology is that only one application ranking service was used (StateOfTheDApps.com). This means that while it is a widely recognised ranking service, several applications may be missing which may provide useful and large data sets for further academic research. This can be addressed by incorporating the results from multiple ranking services simultaneously, which would be analagous to using the Apple App Store, the Google Play Store, and other mobile app stores together in order to assess the size of the mobile app market.

The final three studies (4, 5, and 6) can also be broadly grouped together as they each use transaction data gathered from the Ethereum blockchain. They therefore suffer all of the shortcomings of all player behaviour tracking research [1], the most important of which is that despite the actual behaviours of players of these games being known, no data exists regarding their circumstances which is very important when contextualising their gambling behaviours. Additionally, no demographic data exists describing the populations studied in this thesis, leaving inferential descriptions the current

best option for assessing the player's characteristics. This limitation may be overcome by taking a survey-based approach, seeking to gather demographic and circumstantial data on players and their cryptocurrency address, but this would be enough work to fill a second thesis.

This limitation of the circumstances of the gamblers studied in this thesis being unknown may also be overcome by performing a deeper exploration of their on-chain transactions. Specifically which other platforms they have engaged with, including non-gambling applications, and other transfers of cryptocurrency to/from their addresses. This may also help shed light on the way in which users of decentralised gambling applications buy and sell the currency itself. As described in Section 1.4, not knowing when users purchase an amount of cryptocurrency which is then used for gambling means that it is difficult to express the value of the bets placed in terms which are 'real' to the player. In the studies throughout this thesis it may be that all of the bets are from currency first acquired in 2015, making the bet-time value of each bet very low. It may however be the case that cryptocurrency is bought immediately before the bets are placed, making the real value transacted closer to that of other online casinos. This limitation does not mean that the explorations set out in this thesis are void, but rather that they approach the problem from a cryptocurrency-centric perspective, making further and broader economics-focused work essential to holistically understanding this emerging domain.

A final limitation spanning all of the studies in this thesis is that all of the data studied was derived from the Ethereum blockchain. This means that although findings may apply to decentralised gambling applications implemented using other cryptocurrencies, the exact generalisability of the findings remains unknown, especially when considering networks with different architectures which offer lower transaction confirmation times - and therefore more responsive casino games. This limitation could be overcome by repeating all of the studies presented on new data, although the different languages, features, and configurations, of different cryptocurrency networks makes this process very resource intensive.

10.3 Further Work

All of the studies in this thesis present only the tip of the iceberg in terms of what is now possible in player tracking research given the availability of large scale in-vivo data sets provided by decentralised gambling applications. Unfortunately, existing player behaviour tracking research as a field is not developed enough to robustly identify those at risk of experiencing gambling related harms using transaction data in new domains. However, informative results can still be obtained by taking a comparative approach as in Studies 4, 5, and 6, framing interpretations of results against domains in which academic work has been established.

Having completed this thesis, the most obvious and likely impactful area of future work would be the creation of real time blockchain analytics tools similar to the ranking services described in this work, but which could provide each of the analyses in the studies presented in real time. The pace of innovation in the cryptocurrency and decentralised application space is such that despite their creation as early as 2017, only in this thesis and in limited recent work have decoded transactions to these applications actually been used. Analyses available in real time would be an excellent resource for all of the stakeholders of the work in this thesis, and could be expanded to include multiple cryptocurrencies, multiple applications, and other data sources over time.

A second avenue of useful future work building on the work presented above would be to gather qualitative data sets which could provide broader context to the transaction data freely available on cryptocurrency blockchains. Such qualitative data such as demographic information, income data, and other broader information could be used to compare the profiles of players in a more holistic way, therefore building a more complete picture of potential harms in this emerging domain. Such data would also complement the cluster descriptions in studies like Study 6 in this thesis, providing context for different types of engagement with these applications.

Despite the systematic review revealing the field of player behaviour tracking to be still in its infancy, another exciting avenue of further work is to continue the trend of exploring new behavioural measures which use new and different computations on gambling transaction data. The emerging domain of cryptocurrency gambling via decentralised applications will likely continue

to offer large (and growing) data sets for academic research, upon which new behavioural measures could be tested. As discussed in Section 4.4.4, it is important to choose meaningful outcome variables in such measure-oriented research, but these can be gathered via qualitative methods described above, and may find equal application outside of cryptocurrency research.

Finally, it is clear from the systematic review in Chapter 5 that the behavioural profiling capabilities available to academics are far behind those used by the gambling industry. One way to close this gap would be to take an open source approach to the algorithms used in player behaviour tracking, such that researchers could easily access the code and data used by previous authors. Unfortunately, this is as much a cultural movement as it is an area of future work, but one which I hope will gain support as the field matures.

*“I go, and it is done; the bell invites me.
Hear it not, Duncan, for it is a knell
That summons thee to heaven or to hell.”*

Appendix A

Study 2 Raw Plots

The plots in this section were originally written into Study 2 in Chapter [6](#), but revealed the presence of two large and misclassified applications (Etherpromoswin and CoinGathernator). This spurred the removal of these applications from the study, and a re-run of the visualisations which now appear in the chapter.

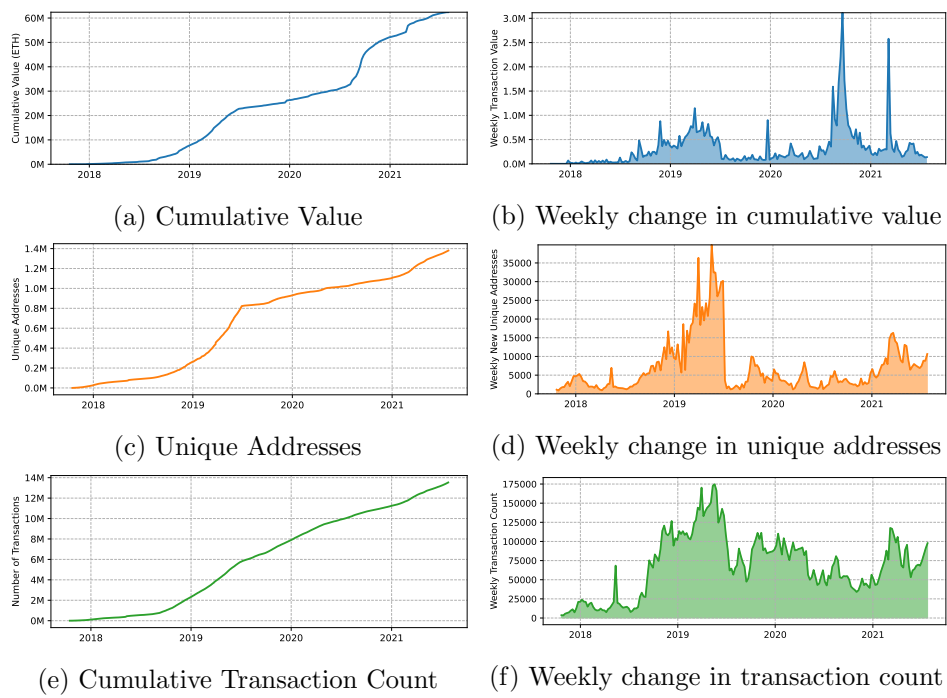





Figure A.1: Usage metrics across all 147 decentralised gambling applications gathered as part of this thesis.



(a) Total Value



(b) Unique Addresses (Player) Counts



(c) Transaction Counts

Figure A.2: Usage metrics of the top 10 decentralised gambling applications in the data sample gathered for this study.

Appendix B

Study 4 Additional Results

The contents of Chapter 7 was published in a PLOS One article [87], but has a limitation that it is now (15-03-2022) somewhat out of date. This appendix presents the results of the same code executed across a more recent version (up to 26-07-2021) of the data set. Additionally, the cutoff of duration less than 4 is applied, so that the results can be *directly* compared with LaBrie et al’s original study [2], rather than the broader comparison provided in Chapter 7. For each of the tables in this section, n=4,328.

Table B.1: Behavioural measure scores across the entire sample.

	Mean	Std	Median	IQR	K-S Score	p-value
duration	101.465	158.528	31.000	108.000	1.000	0.000
frequency_percent	31.038	29.720	20.000	43.182	0.904	0.000
number_of_bets	428.514	1,618.011	70.000	272.000	0.977	0.000
mean_bets_per_day	29.381	46.785	12.000	32.636	0.872	0.000
mean_bet_size	0.907	3.264	0.194	0.522	0.504	0.000
total_wagered	564.125	5,391.429	14.770	101.270	0.733	0.000
net_loss	4.649	85.077	0.199	2.979	0.254	0.000
percent_loss	5.780	58.698	3.091	21.251	0.523	0.000

Table B.2: Spearman's rank correlation coefficient between behavioural measures across all players.

	dur	freq	num	mday	msize	tot	nloss	
dur	-							
freq	-0.8**	-						
num	0.19**	0.29**	-					
mday	-0.05**	0.32**	0.9**	-				
msize	0.03*	0.07**	0.2**	0.18**	-			
tot	0.15**	0.25**	0.82**	0.74**	0.7**	-		
nloss	0.06**	-0.04*	0.13**	0.15**	0.14**	0.18**	-	
ploss	-0.010	-0.1**	-0.15**	-0.1**	-0.07**	-0.14**	0.71**	-

Table B.3: Heavily involved and non-heavily involved split across all players. Total player counts are lower than the study presented in Chapter 7 as although the data covers a longer period of time, players with duration less than 4 (9,241 out of 13,569) have been removed.

	Top 5% ($n = 217$)				Other 95% ($n = 4,111$)			
	Mean	Std	Median	IQR	Mean	Std	Median	IQR
dur	195.645	227.435	111.000	251.000	96.494	152.468	29.000	101.000
freq	40.088	31.751	29.545	49.580	30.560	29.536	20.000	43.348
bets	3,477.724	5,921.406	1,559.000	3,603.000	267.561	629.796	59.000	223.500
mday	102.387	91.875	77.979	90.616	25.527	39.554	10.786	28.361
msize	5.340	7.266	2.333	5.048	0.673	2.711	0.175	0.432
tot	9,389.143	22,339.823	2,805.210	4,209.180	98.295	211.732	12.120	70.260
nloss	58.334	362.748	16.543	169.475	1.815	23.337	0.188	2.502
ploss	0.569	5.200	0.673	4.300	6.055	60.203	3.545	22.863

Bibliography

- [1] B. T. Chagas and J. F. Gomes, “Internet gambling: A critical review of behavioural tracking research,” *Journal of Gambling Issues*, vol. 36, 2017.
- [2] R. A. LaBrie, S. A. Kaplan, D. A. LaPlante, S. E. Nelson, and H. J. Shaffer, “Inside the virtual casino: A prospective longitudinal study of actual internet casino gambling,” *European journal of public health*, vol. 18, no. 4, pp. 410–416, 2008.
- [3] J. Braverman and H. J. Shaffer, “How do gamblers start gambling: Identifying behavioural markers for high-risk internet gambling,” *The European Journal of Public Health*, vol. 22, no. 2, pp. 273–278, 2010.
- [4] R. K. L. Nielsen and P. Grabarczyk, “Are loot boxes gambling? random reward mechanisms in video games,” *Transactions of the Digital Games Research Association*, vol. 4, no. 3, 2019.
- [5] S. M. Gainsbury and A. Blaszczynski, “How blockchain and cryptocurrency technology could revolutionize online gambling,” *Gaming Law Review*, vol. 21, no. 7, pp. 482–492, 2017.
- [6] S. I. Millar, “Cryptocurrency expands online gambling,” *Gaming Law Review*, vol. 22, no. 3, pp. 174–174, 2018.
- [7] S. Lawn, C. Oster, B. Riley, D. Smith, M. Baigent, and M. Rahamathulla, “A literature review and gap analysis of emerging technologies and new trends in gambling,” *International journal of environmental research and public health*, vol. 17, no. 3, p. 744, 2020.
- [8] S. R. Currie, D. C. Hodgins, J. Wang, N. El-Guebaly, H. Wynne, and S. Chen, “Risk of harm among gamblers in the general population as a function of level of participation in gambling activities,” *Addiction*, vol. 101, no. 4, pp. 570–580, 2006.
- [9] O. J. Scholten, N. G. J. Hughes, S. Deterding, A. Drachen, J. A. Walker, and D. Zendle, “Ethereum crypto-games: Mechanics, prevalence, and gambling similarities,” in *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, pp. 379–389, 2019.
- [10] E. Barker, “Guideline for using cryptographic standards in the federal government: Cryptographic mechanisms,” *NIST Special Publication*, 2020.
- [11] J. Katz and Y. Lindell, *Introduction to modern cryptography*. CRC press, 2020.
- [12] N. I. o. S. Information Technology Laboratory and Technology, “Secure hash standard (shs),” *Federal Information Processing Standards Publication*, 2015.
- [13] I. Dinur, O. Dunkelman, and A. Shamir, “New attacks on keccak-224 and keccak-256,” in *International Workshop on Fast Software Encryption*, pp. 442–461, Springer, 2012.
- [14] J. Deepakumara, H. M. Heys, and R. Venkatesan, “Fpga implementation of md5 hash algorithm,” in *Canadian Conference on Electrical and Computer Engineering 2001. Conference Proceedings (Cat. No. 01TH8555)*, vol. 2, pp. 919–924, IEEE, 2001.

- [15] J. P. Conley *et al.*, *Encryption, Hashing, PPK, and Blockchain: A Simple Introduction*. Vanderbilt University, Department of Economics, 2019.
- [16] S. Goldwasser, S. Micali, and R. L. Rivest, “A digital signature scheme secure against adaptive chosen-message attacks,” *SIAM Journal on computing*, vol. 17, no. 2, pp. 281–308, 1988.
- [17] R. L. Rivest, A. Shamir, and L. Adleman, “A method for obtaining digital signatures and public-key cryptosystems,” *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.
- [18] T. ElGamal, “A public key cryptosystem and a signature scheme based on discrete logarithms,” *IEEE transactions on information theory*, vol. 31, no. 4, pp. 469–472, 1985.
- [19] J. Hoffstein, N. Howgrave-Graham, J. Pipher, J. H. Silverman, and W. Whyte, “NtruSign: Digital signatures using the ntru lattice,” in *Cryptographers’ track at the RSA conference*, pp. 122–140, Springer, 2003.
- [20] M. Pilkington, “Blockchain technology: principles and applications,” in *Research handbook on digital transformations*, Edward Elgar Publishing, 2016.
- [21] A. D. Kshemkalyani and M. Singhal, *Distributed computing: principles, algorithms, and systems*. Cambridge University Press, 2011.
- [22] V. S. Pande, I. Baker, J. Chapman, S. P. Elmer, S. Khaliq, S. M. Larson, Y. M. Rhee, M. R. Shirts, C. D. Snow, E. J. Sorin, *et al.*, “Atomistic protein folding simulations on the submillisecond time scale using worldwide distributed computing,” *Biopolymers: Original Research on Biomolecules*, vol. 68, no. 1, pp. 91–109, 2003.
- [23] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong, “Freenet: A distributed anonymous information storage and retrieval system,” in *Designing privacy enhancing technologies*, pp. 46–66, Springer, 2001.
- [24] I. Clarke, O. Sandberg, M. Toseland, and V. Verendel, “Private communication through a network of trusted connections: The dark freenet,” *Network*, 2010.
- [25] B. Zantout, R. Haraty, *et al.*, “I2p data communication system,” in *Proceedings of ICN*, pp. 401–409, Citeseer, 2011.
- [26] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The second-generation onion router,” tech. rep., Naval Research Lab Washington DC, 2004.
- [27] V. Buterin *et al.*, “A next-generation smart contract and decentralized application platform,”
- [28] W. Wang, D. T. Hoang, P. Hu, Z. Xiong, D. Niyato, P. Wang, Y. Wen, and D. I. Kim, “A survey on consensus mechanisms and mining strategy management in blockchain networks,” *IEEE Access*, vol. 7, pp. 22328–22370, 2019.
- [29] Y. Yuan and F.-Y. Wang, “Blockchain and cryptocurrencies: Model, techniques, and applications,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 48, no. 9, pp. 1421–1428, 2018.
- [30] C. T. Nguyen, D. T. Hoang, D. N. Nguyen, D. Niyato, H. T. Nguyen, and E. Dutkiewicz, “Proof-of-stake consensus mechanisms for future blockchain networks: fundamentals, applications and opportunities,” *IEEE Access*, vol. 7, pp. 85727–85745, 2019.
- [31] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” tech. rep., Manubot, 2008.
- [32] D. Chaum, “Blind signatures for untraceable payments,” in *Advances in cryptology*, pp. 199–203, Springer, 1983.
- [33] W. Dai, “B-money-an anonymous, distributed electronic cash system,” 1998.

- [34] T. H. Kim, "A study of digital currency cryptography for business marketing and finance security," *Asia-pacific Journal of Multimedia Services Convergent with Art, Humanities, and Sociology*, vol. 6, no. 1, pp. 365–376, 2016.
- [35] N. Szabo, "Smart contracts: building blocks for digital markets," *EXTROPY: The Journal of Transhumanist Thought*, (16), vol. 18, no. 2, 1996.
- [36] N. Szabo, "Formalizing and securing relationships on public networks," *First monday*, 1997.
- [37] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," tech. rep., Manubot, 2019.
- [38] V. Lehdonvirta and E. Castronova, *Virtual economies: Design and analysis*. MIT Press, 2014.
- [39] G. Hileman, "From bitcoin to the brixton pound: history and prospects for alternative currencies," in *International Conference on Financial Cryptography and Data Security*, pp. 163–165, Springer, 2014.
- [40] A. Serada, T. Sihvonen, and J. T. Harviainen, "Cryptokitties and the new ludic economy: how blockchain introduces value, ownership, and scarcity in digital gaming," *Games and Culture*, p. 1555412019898305, 2020.
- [41] P. Wang, X. Li, D. Shen, and W. Zhang, "How does economic policy uncertainty affect the bitcoin market?," *Research in International Business and Finance*, vol. 53, p. 101234, 2020.
- [42] H.-P. Cheng and K.-C. Yen, "The relationship between the economic policy uncertainty and the cryptocurrency market," *Finance Research Letters*, vol. 35, p. 101308, 2020.
- [43] O. Angela and Y. Sun, "Factors affecting cryptocurrency prices: Evidence from ethereum," in *2020 International Conference on Information Management and Technology (ICIMTech)*, pp. 318–323, IEEE, 2020.
- [44] S. Bhambhwani, S. Delikouras, G. M. Korniotis, *et al.*, *Do fundamentals drive cryptocurrency prices?* Centre for Economic Policy Research, 2019.
- [45] N. Gandal, J. Hamrick, T. Moore, and T. Oberman, "Price manipulation in the bitcoin ecosystem," *Journal of Monetary Economics*, vol. 95, pp. 86–96, 2018.
- [46] T. Li, D. Shin, and B. Wang, "Cryptocurrency pump-and-dump schemes," *Available at SSRN 3267041*, 2021.
- [47] L. Nizzoli, S. Tardelli, M. Avvenuti, S. Cresci, M. Tesconi, and E. Ferrara, "Charting the landscape of online cryptocurrency manipulation," *IEEE Access*, vol. 8, pp. 113230–113245, 2020.
- [48] F. Stajano and R. Clayton, "Cyberdice: peer-to-peer gambling in the presence of cheaters," in *International Workshop on Security Protocols*, pp. 54–70, Springer, 2008.
- [49] T. Min, H. Wang, Y. Guo, and W. Cai, "Blockchain games: A survey," in *2019 IEEE Conference on Games (CoG)*, pp. 1–8, IEEE, 2019.
- [50] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, no. 2014, pp. 1–32, 2014.
- [51] P. McCorry, S. F. Shahandashti, and F. Hao, "A smart contract for boardroom voting with maximum voter privacy," in *International conference on financial cryptography and data security*, pp. 357–375, Springer, 2017.
- [52] M. Auer and M. D. Griffiths, "Reasons for gambling and problem gambling among norwegian horse bettors: A real-world study utilizing combining survey data and behavioral player data," *International Journal of Mental Health and Addiction*, pp. 1–16, 2021.

- [53] I. Ukhov, J. Bjurgert, M. Auer, and M. D. Griffiths, "Online problem gambling: a comparison of casino players and sports bettors via predictive modeling using behavioral tracking data," *Journal of Gambling Studies*, pp. 1–21, 2020.
- [54] D. G. Schwartz, "Roll the bones: The history of gambling," 2013.
- [55] M. D. Griffiths, "Is the buying of loot boxes in video games a form of gambling or gaming?," *Gaming Law Review*, vol. 22, no. 1, pp. 52–54, 2018.
- [56] M. Griffiths, *Adolescent gambling*. Psychology Press, 1995.
- [57] S. M. Gainsbury, N. Hing, P. H. Delfabbro, and D. L. King, "A taxonomy of gambling and casino games via social media and online technologies," *International Gambling Studies*, vol. 14, no. 2, pp. 196–213, 2014.
- [58] S. M. Gainsbury, A. M. Russell, D. L. King, P. Delfabbro, and N. Hing, "Migration from social casino games to gambling: Motivations and characteristics of gamers who gamble," *Computers in Human Behavior*, vol. 63, pp. 59–67, 2016.
- [59] K. Dunphy, "Following suit with the second circuit: Defining gambling in the illegal gambling business act," *Brook. L. Rev.*, vol. 79, p. 1295, 2013.
- [60] J. M. Roberts and B. Sutton-Smith, "Cross-cultural correlates of games of chance," *Behavior Science Notes*, vol. 1, no. 3, pp. 131–144, 1966.
- [61] R. A. LaBrie, D. A. LaPlante, S. E. Nelson, A. Schumann, and H. J. Shaffer, "Assessing the playing field: A prospective longitudinal study of internet sports gambling behavior," *Journal of Gambling studies*, vol. 23, no. 3, pp. 347–362, 2007.
- [62] M. Gusenbauer and N. R. Haddaway, "Which academic search systems are suitable for systematic reviews or meta-analyses? evaluating retrieval qualities of google scholar, pubmed, and 26 other resources," *Research synthesis methods*, vol. 11, no. 2, pp. 181–217, 2020.
- [63] I. Fiedler, "Gamblers' habits: Empirical evidence on the behavior of regulars, newcomers and dropouts," *Journal of gambling studies*, vol. 29, no. 2, pp. 289–309, 2013.
- [64] M. Auer and M. D. Griffiths, "Self-reported losses versus actual losses in online gambling: An empirical study," *Journal of gambling studies*, vol. 33, no. 3, pp. 795–806, 2017.
- [65] I. Fiedler, "The gambling habits of online poker players," *The Journal of Gambling Business and Economics*, 2011.
- [66] A. Luquiens, D. Vendryes, H.-J. Aubin, A. Benyamina, S. Gaiffas, and E. Bacry, "Description and assessment of trustability of motives for self-exclusion reported by online poker gamblers in a cohort using account-based gambling data," *BMJ open*, vol. 8, no. 12, p. e022541, 2018.
- [67] M. Auer and M. D. Griffiths, "An empirical investigation of theoretical loss and gambling intensity," *Journal of Gambling Studies*, vol. 30, no. 4, pp. 879–887, 2014.
- [68] J. Haeusler, "Follow the money: using payment behaviour as predictor for future self-exclusion," *International Gambling Studies*, vol. 16, no. 2, pp. 246–262, 2016.
- [69] J. Braverman, D. A. LaPlante, S. E. Nelson, and H. J. Shaffer, "Using cross-game behavioral markers for early identification of high-risk internet gamblers," *Psychology of Addictive Behaviors*, vol. 27, no. 3, p. 868, 2013.
- [70] D. Sagoe, S. Pallesen, M. D. Griffiths, R. A. Mentzoni, and T. Leino, "Does individual gambling behavior vary across gambling venues with differing numbers of terminals? an empirical real-world study using player account data," *Frontiers in psychology*, vol. 9, p. 158, 2018.

- [71] M. Catania and M. D. Griffiths, "Understanding online voluntary self-exclusion in gambling: An empirical study using account-based behavioral tracking data," *International Journal of Environmental Research and Public Health*, vol. 18, no. 4, p. 2000, 2021.
- [72] C. Percy, M. França, S. Dragičević, and A. d'Avila Garcez, "Predicting online gambling self-exclusion: an analysis of the performance of supervised machine learning models," *International Gambling Studies*, vol. 16, no. 2, pp. 193–210, 2016.
- [73] J. Jonsson, D. C. Hodgins, I. Munck, and P. Carlbring, "Reaching out to big losers: A randomized controlled trial of brief motivational contact providing gambling expenditure feedback.," *Psychology of Addictive Behaviors*, vol. 33, no. 3, p. 179, 2019.
- [74] M. J. Wohl, C. G. Davis, and S. J. Hollingshead, "How much have you won or lost? personalized behavioral feedback about gambling expenditures regulates play," *Computers in Human Behavior*, vol. 70, pp. 437–445, 2017.
- [75] M. Auer and M. D. Griffiths, "Voluntary limit setting and player choice in most intense online gamblers: An empirical study of gambling behaviour," *Journal of Gambling Studies*, vol. 29, no. 4, pp. 647–660, 2013.
- [76] M. Auer and M. D. Griffiths, "Gambling before and during the covid-19 pandemic among online casino gamblers: an empirical study using behavioral tracking data," *International Journal of Mental Health and Addiction*, pp. 1–11, 2021.
- [77] M. Auer and M. D. Griffiths, "The use of personalized messages on wagering behavior of swedish online gamblers: An empirical study," *Computers in Human Behavior*, vol. 110, p. 106402, 2020.
- [78] G. Challet-Bouju, J.-B. Hardouin, E. Thiabaud, A. Saillard, Y. Donnio, M. Grall-Bronnec, and B. Perrot, "Modeling early gambling behavior using indicators from online lottery gambling tracking data: Longitudinal analysis," *Journal of Medical Internet Research*, vol. 22, no. 8, p. e17675, 2020.
- [79] M. Auer, N. Hopfgartner, and M. D. Griffiths, "The effects of a mandatory play break on subsequent gambling among norwegian video lottery terminal players," *Journal of behavioral addictions*, vol. 8, no. 3, pp. 522–529, 2019.
- [80] D. Forsström, H. Hesser, and P. Carlbring, "Usage of a responsible gambling tool: a descriptive analysis and latent class analysis of user behavior," *Journal of Gambling Studies*, vol. 32, no. 3, pp. 889–904, 2016.
- [81] K. S. Philander, "Identifying high-risk online gamblers: a comparison of data mining procedures," *International Gambling Studies*, vol. 14, no. 1, pp. 53–63, 2014.
- [82] S. Sarkar, T. Weyde, A. Garcez, G. G. Slabaugh, S. Dragicevic, and C. Percy, "Accuracy and interpretability trade-offs in machine learning applied to safer gambling," in *CEUR Workshop Proceedings*, vol. 1773, CEUR Workshop Proceedings, 2016.
- [83] S. M. Gainsbury and A. Russell, "Betting patterns for sports and races: A longitudinal analysis of online wagering in australia," *Journal of Gambling Studies*, vol. 31, no. 1, pp. 17–32, 2015.
- [84] M. M. Auer and M. D. Griffiths, "Personalized behavioral feedback for online gamblers: A real world empirical study," *Frontiers in Psychology*, vol. 7, p. 1875, 2016.
- [85] M. M. Auer and M. D. Griffiths, "The use of personalized behavioral feedback for online gamblers: an empirical study," *Frontiers in psychology*, vol. 6, p. 1406, 2015.
- [86] M. Auer, N. Hopfgartner, and M. D. Griffiths, "The effect of loss-limit reminders on gambling behavior: A real-world study of norwegian gamblers," *Journal of Behavioral Addictions*, vol. 7, no. 4, pp. 1056–1067, 2018.

- [87] O. J. Scholten, D. Zendle, and J. A. Walker, "Inside the decentralised casino: A longitudinal study of actual cryptocurrency gambling transactions," *PloS one*, vol. 15, no. 10, p. e0240693, 2020.
- [88] D. A. LaPlante, A. Schumann, R. A. LaBrie, and H. J. Shaffer, "Population trends in internet sports gambling," *Computers in Human Behavior*, vol. 24, no. 5, pp. 2399–2414, 2008.
- [89] M. M. Auer and M. D. Griffiths, "Testing normative and self-appraisal feedback in an online slot-machine pop-up in a real-world setting," *Frontiers in psychology*, vol. 6, p. 339, 2015.
- [90] T. Leino, T. Torsheim, A. Blaszczynski, M. Griffiths, R. Mentzoni, S. Pallesen, and H. Molde, "The relationship between structural game characteristics and gambling behavior: A population-level study," *Journal of gambling studies*, vol. 31, no. 4, pp. 1297–1315, 2015.
- [91] H. M. Gray, G. K. Jónsson, D. A. LaPlante, and H. J. Shaffer, "Expanding the study of internet gambling behavior: Trends within the icelandic lottery and sportsbetting platform," *Journal of Gambling Studies*, vol. 31, no. 2, pp. 483–499, 2015.
- [92] Z. Xuan and H. Shaffer, "How do gamblers end gambling: Longitudinal analysis of internet gambling behaviors prior to account closure due to gambling related problems," *Journal of Gambling Studies*, vol. 25, no. 2, pp. 239–252, 2009.
- [93] S. Dragicevic, G. Tsogas, and A. Kudic, "Analysis of casino online gambling data in relation to behavioural risk markers for high-risk gambling and player protection," *International Gambling Studies*, vol. 11, no. 3, pp. 377–391, 2011.
- [94] R. LaBrie and H. J. Shaffer, "Identifying behavioral markers of disordered internet sports gambling," *Addiction Research & Theory*, vol. 19, no. 1, pp. 56–65, 2011.
- [95] H. M. Gray, D. A. LaPlante, and H. J. Shaffer, "Behavioral characteristics of internet gamblers who trigger corporate responsible gambling interventions," *Psychology of Addictive Behaviors*, vol. 26, no. 3, p. 527, 2012.
- [96] J. Braverman, R. A. LaBrie, and H. J. Shaffer, "A taxometric analysis of actual internet sports gambling behavior," *Psychological Assessment*, vol. 23, no. 1, p. 234, 2011.
- [97] T. C. Edson, M. A. Tom, K. S. Philander, E. R. Louderback, and D. A. LaPlante, "A large-scale prospective study of big wins and their relationship with future financial and time involvement in actual daily fantasy sports contests," *Psychology of Addictive Behaviors*, 2021.
- [98] S. E. Nelson, D. A. LaPlante, A. J. Peller, A. Schumann, R. A. LaBrie, and H. J. Shaffer, "Real limits in the virtual world: Self-limiting behavior of internet gamblers," *Journal of gambling Studies*, vol. 24, no. 4, pp. 463–477, 2008.
- [99] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.
- [100] D. A. LaPlante, J. H. Kleschinsky, R. A. LaBrie, S. E. Nelson, and H. J. Shaffer, "Sitting at the virtual poker table: A prospective epidemiological study of actual internet poker gambling behavior," *Computers in Human Behavior*, vol. 25, no. 3, pp. 711–717, 2009.
- [101] S. Finkenwirth, K. MacDonald, X. Deng, T. Lesch, and L. Clark, "Using machine learning to predict self-exclusion status in online gamblers on the playnow. com platform in british columbia," *International Gambling Studies*, pp. 1–18, 2020.
- [102] S. Dragicevic, C. Percy, A. Kudic, and J. Parke, "A descriptive analysis of demographic and behavioral data from internet gamblers and those who self-exclude from

- online gambling platforms,” *Journal of Gambling Studies*, vol. 31, no. 1, pp. 105–132, 2015.
- [103] S. Gainsbury, S. Sadeque, D. Mizerski, A. Blaszczynski, *et al.*, “Wagering in australia: A retrospective behavioural analysis of betting patterns based on player account data,” *Journal of Gambling Business and Economics*, vol. 6, no. 2, p. 50, 2012.
- [104] A. Broda, D. A. LaPlante, S. E. Nelson, R. A. LaBrie, L. B. Bosworth, and H. J. Shaffer, “Virtual harm reduction efforts for internet gambling: effects of deposit limits on actual internet sports gambling behavior,” *Harm reduction journal*, vol. 5, no. 1, p. 27, 2008.
- [105] J. Braverman and H. J. Shaffer, “How do gamblers start gambling: Identifying behavioural markers for high-risk internet gambling,” *The European Journal of Public Health*, vol. 22, no. 2, pp. 273–278, 2010.
- [106] N. Adami, S. Benini, A. Boschetti, L. Canini, F. Maione, and M. Temporin, “Markers of unsustainable gambling for early detection of at-risk online gamblers,” *International Gambling Studies*, vol. 13, no. 2, pp. 188–204, 2013.
- [107] X. Han, L. Wang, S. Xu, D. Zhao, and G. Liu, “Recognizing roles of online illegal gambling participants: an ensemble learning approach,” *Computers & Security*, vol. 87, p. 101588, 2019.
- [108] X. Ma, S. H. Kim, and S. S. Kim, “Online gambling behavior: The impacts of cumulative outcomes, recent outcomes, and prior use,” *Information Systems Research*, vol. 25, no. 3, pp. 511–527, 2014.
- [109] M. Auer and M. D. Griffiths, “Predicting limit-setting behavior of gamblers using machine learning algorithms: A real-world study of norwegian gamblers using account data,” *International Journal of Mental Health and Addiction*, pp. 1–18, 2019.
- [110] T. Kainulainen, “Does losing on a previous betting day predict how long it takes to return to the next session of online horse race betting?,” *Journal of Gambling Studies*, vol. 37, no. 2, pp. 609–622, 2021.
- [111] J. Braverman, M. Tom, and H. J. Shaffer, “Tilting at windmills: A comment on auer and griffiths,” *Journal of gambling studies*, vol. 31, no. 2, pp. 359–366, 2015.
- [112] M. Auer and M. D. Griffiths, “Theoretical loss and gambling intensity (revisited): A response to braverman et al.(2013),” *Journal of gambling studies*, vol. 31, no. 3, pp. 921–931, 2015.
- [113] J. Meng and F. Fu, “Understanding gambling behaviour and risk attitudes using cryptocurrency-based casino blockchain data,” *Royal Society open science*, vol. 7, no. 10, p. 201446, 2020.
- [114] G. Tian, Q. Wang, Y. Zhao, L. Guo, Z. Sun, and L. Lv, “Smart contract classification with a bi-lstm based approach,” *IEEE Access*, vol. 8, pp. 43806–43816, 2020.
- [115] S. Sezer, C. Eyhoff, W. Prinz, and T. Rose, “Exploiting smart contract bytecode for classification on ethereum,” in *PoEM Workshops*, pp. 11–22, 2020.
- [116] C. Shi, Y. Xiang, R. R. M. Doss, J. Yu, K. Sood, and L. Gao, “A bytecode-based approach for smart contract classification,” *arXiv preprint arXiv:2106.15497*, 2021.
- [117] A. Koivula, P. Räsänen, I. Koiranen, and T. Keipi, “Online gambling activity in finland 2006–2016,” in *International Conference on Human-Computer Interaction*, pp. 569–583, Springer, 2020.
- [118] E. Saiedi, A. Broström, and F. Ruiz, “Global drivers of cryptocurrency infrastructure adoption,” *Small Business Economics*, vol. 57, no. 1, pp. 353–406, 2021.
- [119] S. Kethineni and Y. Cao, “The rise in popularity of cryptocurrency and associated criminal activity,” *International Criminal Justice Review*, vol. 30, no. 3, pp. 325–344, 2020.

- [120] C. Chatfield and H. Xing, *The analysis of time series: an introduction with R*. Chapman and hall/CRC, 2019.
- [121] L. Ante, “Non-fungible token (nft) markets on the ethereum blockchain: Temporal development, cointegration and interrelations,” *Available at SSRN 3904683*, 2021.
- [122] M. Dowling, “Is non-fungible token pricing driven by cryptocurrencies?,” *Finance Research Letters*, p. 102097, 2021.
- [123] S. Gainsbury, “Player account-based gambling: Potentials for behaviour-based research methodologies,” *International Gambling Studies*, vol. 11, no. 2, pp. 153–171, 2011.
- [124] D. Johnson, A. Menezes, and S. Vanstone, “The elliptic curve digital signature algorithm (ecdsa),” *International journal of information security*, vol. 1, no. 1, pp. 36–63, 2001.
- [125] S. E. Nelson, T. C. Edson, E. R. Louderback, M. A. Tom, A. Grossman, and D. A. LaPlante, “Changes to the playing field: A contemporary study of actual european online sports betting,” *Journal of Behavioral Addictions*, 2021.
- [126] A. Drummond and J. D. Sauer, “Video game loot boxes are psychologically akin to gambling,” *Nature human behaviour*, vol. 2, no. 8, pp. 530–532, 2018.
- [127] D. Zendle and P. Cairns, “Video game loot boxes are linked to problem gambling: Results of a large-scale survey,” *PloS one*, vol. 13, no. 11, p. e0206767, 2018.
- [128] D. Zendle and P. Cairns, “Loot boxes are again linked to problem gambling: Results of a replication study,” *PloS one*, vol. 14, no. 3, p. e0213194, 2019.
- [129] D. Zendle, P. Cairns, H. Barnett, and C. McCall, “Paying for loot boxes is linked to problem gambling, regardless of specific features like cash-out and pay-to-win,” *Computers in Human Behavior*, vol. 102, pp. 181–191, 2020.
- [130] X.-J. Jiang and X. F. Liu, “Cryptokitties transaction network analysis: The rise and fall of the first blockchain game mania,” *Frontiers in Physics*, vol. 9, p. 57, 2021.
- [131] A. Likas, N. Vlassis, and J. J. Verbeek, “The global k-means clustering algorithm,” *Pattern recognition*, vol. 36, no. 2, pp. 451–461, 2003.
- [132] D. Arthur and S. Vassilvitskii, “How slow is the k-means method?,” in *Proceedings of the twenty-second annual symposium on Computational geometry*, pp. 144–153, 2006.
- [133] T. M. Kodinariya and P. R. Makwana, “Review on determining number of cluster in k-means clustering,” *International Journal*, vol. 1, no. 6, pp. 90–95, 2013.
- [134] F. Steinmetz, M. von Meduna, L. Ante, and I. Fiedler, “Ownership, uses and perceptions of cryptocurrency: Results from a population survey,” *Technological Forecasting and Social Change*, vol. 173, p. 121073, 2021.
- [135] F. Edition *et al.*, “Diagnostic and statistical manual of mental disorders,” *Am Psychiatric Assoc*, vol. 21, 2013.
- [136] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [137] N. D. Schüll, *Addiction by design*. Princeton University Press, 2012.