Health, Income, and Employment: Empirical Insights into Economic, Behavioural, and Genetic Factors.

Aharon Katz

Doctor of Philosophy

University of York

Department of Economics and Related Studies

June 2025

Abstract

This thesis examines the interrelationships between health, income, and employment using empirical methods that integrate economic, behavioural, and genetic perspectives. It consists of three chapters, each addressing a key question at the intersection of health and labour market outcomes, with a strong focus on causal inference.

The first chapter investigates the impact of mental health on early retirement among individuals aged 50–64 in Australia, using data from the Household, Income and Labour Dynamics in Australia (HILDA) survey. Linear models, survival analysis, and instrumental variable methods are applied to address endogeneity, with mental health instrumented using the death of a family member in the previous 12 months. The results show a positive and statistically significant association between poor mental health and the probability of early retirement, highlighting the importance of mental health as a determinant of labour market exit.

The second chapter examines the relationship between income and health in Israel, distinguishing between absolute and relative income effects. Relative income is measured both at the population level and within major religious groups. The analysis applies OLS, Probit models with fixed and random effects, Hausman–Taylor estimation, and semiparametric techniques. The findings indicate a significant positive relationship between absolute income and health, with stronger marginal effects observed for relative income. The results also reveal heterogeneity across religious groups, suggesting that social and group dynamics shape the income–health relationship.

The third chapter explores the causal effect of obesity on employment in the United Kingdom using UK Biobank data. Mendelian Randomisation is applied, using a genetic risk score for body mass index as an instrumental variable for obesity. To account for unobserved heterogeneity, marginal treatment effect and person-centered treatment effect methods are applied. The findings show a significant negative effect of obesity on employment probability, with PeT estimates indicating stronger effects than conventional OLS and broadly consistent with MR results.

This thesis is dedicated to the memory of my beloved father and to my devoted mother.

Acknowledgements

First and foremost, I dedicate this thesis to the memory of my beloved father, Avraham Katz (RIP). Though he is no longer with us, his enduring spirit and unwavering emotional support guided me throughout this Ph.D. journey and continue to inspire me every day. I am also deeply grateful to my mother, Ili Katz, whose steadfast love, encouragement, and strength have been a constant source of support. Beginning my Ph.D. just prior to the COVID-19 pandemic brought unique challenges, navigating isolation during multiple lockdowns while being far from my mother and close friends was especially difficult and would not be possible without their support.

I would also like to express my deepest gratitude to my supervisors, Professor Andrew M. Jones and Professor Nigel Rice. I am truly thankful for their generous support throughout my Ph.D. Their patience in sharing their exceptional knowledge and expertise, their meticulous feedback at every stage, and their encouragement to engage critically and reflectively with my work have all been invaluable. Their commitment and passion for research have deeply inspired me. Their insights, advanced knowledge, and experience were a continuous source of motivation and played a crucial role in shaping the three chapters of this thesis.

My sincere thanks go as well to Dr James Lomas, my thesis advisory panel member, whose thoughtful comments and suggestions significantly enhanced this work. I would also like to acknowledge the HEDG cluster for many fruitful discussions and the intellectually stimulating environment. A heartfelt thank you to Dr Michael Shallcross, Maigen Savory, and Vin McDermott for their ongoing support throughout my time at York.

Last but not least, I extend my heartfelt thanks to my dear friends for making this journey a truly warm and memorable experience. In particular, I am deeply grateful to Orel Barda, David Dvora, Alexandra Quinones Nunura, Adam Stirling, and Boen Liang. Your companionship, especially during the isolating lockdown periods, kept me grounded and uplifted. The shared coffees and teas, meaningful conversations, and moments of laughter offered much-needed energy and joy throughout this journey.

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. All sources are acknowledged as references. I have orally presented earlier versions of the three papers at seminars at the Department of Economics and Related Studies at the University of York. Chapter 3 was discussed at the 13th Annual Conference of the American Society of Health Economics (ASHEcon). Chapter 2 and Chapter 3 will be presented at the 16th International Health Economics Association (IHEA) Congress. Chapter 1 of my thesis was published in the *Australian Journal of Labour Economics*, Volume 28, Number 1 (2025).

Aharon Katz
June 2025

"Not everything that counts can be counted, and not everything that can be counted counts."
-Albert Einstein-

Table of contents

List of tables

List of figures

Introduction

The intricate relationship between health and economic activity has become increasingly salient in recent years, especially in light of the COVID-19 pandemic and the accelerating demographic transitions associated with population ageing. These global phenomena have heightened the urgency to understand how health influences individual and collective economic outcomes, particularly in relation to labour supply, productivity, and income. As governments grapple with the dual challenge of strained public health systems and the sustainability of welfare and pension programs, robust empirical evidence on the socioeconomic consequences of health disparities is critical for designing effective policy interventions.

Although the relationship between income and health is well established, most research has focused on the effects of absolute income, that is, how increases in financial resources improve health outcomes. However, less attention has been given to the relative income hypothesis, which emphasises the role of an individual's income position relative to others within their social reference group. A growing body of evidence suggests that relative income can independently affect physical and mental health through psychosocial pathways, such as stress and social comparison (Jones and Wildman, 2008; Mangyo and Park, 2011; Marmot et al., 1984; Subramanian and Kawachi, 2004). Overlooking this dimension risks missing a critical determinant of health inequality. Incorporating both relative and absolute income perspectives is therefore essential for understanding the full scope of socioeconomic influences on health and their downstream implications for economic behaviour.

Obesity has also emerged as a pressing public health and economic concern, with rising prevalence across all demographic groups. It is a key contributor to chronic diseases and has been linked to diminished labour market outcomes, including lower employment probabilities, reduced earnings, and increased absenteeism (Busetta et al., 2020; Cawley, 2004; Li et al., 2022; Morris, 2007). These impacts are shaped not only by direct health-related limitations, but also by broader social factors such as employer discrimination and patterns of self-selection into or out of employment based on characteristics correlated with obesity (e.g., mobility issues, or personal dissatisfaction). The economic costs are substantial, with estimates suggesting that obesity-related healthcare and productivity losses exceed £80 billion annually in the UK alone (Dall et al., 2024; Diabetes UK, 2024). Despite the scale of the issue, identifying causal effects remains empirically challenging due to potential endogeneity. This highlights the need for robust methodological approaches, such as mendelian Randomisation and person-centered treatment effects, to disentangle correlation from causation and inform targeted policy interventions.

At the same time, the demographic shift towards older populations, driven by increased longevity and declining fertility, has placed additional stress on labour markets and public finances (Bloom et al., 2015; Mitra et al., 2020). Rising old-age dependency ratios threaten to reduce labour force participation and amplify demands on pension and healthcare systems. In this context, it is vital to understand how both physical and mental health conditions influence key labour market behaviours such as, early retirement.

This thesis addresses these questions by examining the causal effects of various health dimensions on labour market outcomes, using rich longitudinal data from Australia (The Household, Income and Labour Dynamics in Australia (HILDA)), Israel (The Central Bureau of Statistics (CBS)), and the United Kingdom (UK Biobank). It comprises three empirical chapters, each focusing on a distinct health aspect; mental health, obesity, and general health status, and applying advanced econometric strategies to identify causal pathways. Although set in diverse institutional and demographic environments, the chapters are unified by a shared goal: to deepen our understanding of the health–labour nexus and to provide evidence that can inform effective, data-driven policy.

The first chapter explores the role of mental health in shaping early retirement decisions among older adults, with an analysis based on data from Australia (HILDA). It is an increasingly important issue as ageing populations prompt a reevaluation of traditional retirement norms. Mental illness has emerged as a leading cause of work absence and productivity loss, with estimates suggesting that its economic burden in Australia alone exceeds AU$60 billion annually (Lee et al., 2017; Neil et al., 2014). Despite its growing prevalence, the impact of mental health on labour force participation, particularly in the context of early retirement, remains underexplored relative to physical health. Using discrete-time survival analysis and instrumental variable (IV) techniques, the study addresses endogeneity concerns inherent in self-reported mental health and investigates gender-specific heterogeneities. The findings underscore the significant role of mental health in shaping labour market exit decisions, with implications for disability policy and retirement planning.

The second chapter examines the intersection of health, religion, and income inequality in Israel. The analysis focuses on the relative income hypothesis and investigates how one's position in the income distribution, rather than absolute income alone, affects health outcomes. This study uniquely incorporates religious affiliation (Judaism, Islam, and Christianity) as a potential moderator in the income-health relationship, highlighting the socio-cultural factors that may influence health behaviours and access to resources. Leveraging nationally representative panel data, the analysis applies semiparametric and parametric estimation strategies and constructs a binary health index to mitigate potential measurement error in self-reported health. The findings reveal important interactions between health, income rank across the income distribution, and religious identity, suggesting that policies aimed at reducing income inequality must be sensitive to cultural heterogeneity.

The final empirical chapter applies a Mendelian Randomisation (MR) approach using genetic variants as instrumental variables for body mass index (BMI), drawing on biomarker data from the UK Biobank. This method provides a powerful quasi-experimental framework to address longstanding challenges of reverse causality and omitted variable bias in estimating the impact of obesity on employment and earnings. By exploiting the exogeneity of genetic predispositions, the analysis offers robust causal estimates of the economic costs of excess body weight. The chapter also illustrates the broader potential of integrating genetic data into economic research to improve causal inference and to capture latent health conditions that may not be observed through conventional survey measures.

Across all three chapters, this thesis makes methodological contributions by rigorously addressing issues of endogeneity, selection bias, and unobserved heterogeneity, common pitfalls in applied health economics. Techniques such as instrumental variable estimation (e.g., two-stage residual inclusion (2SRI) and MR) and survival analysis are applied across diverse empirical contexts to ensure the internal validity of findings. Subgroup analyses by gender, age, and religious affiliation further enhance the external validity and policy relevance of the results, offering insights into how health-related labour market outcomes differ across population subgroups.

Taken together, the thesis contributes to the growing literature on the economics of health and labour by producing internationally comparative evidence on the causal effects of health on employment outcomes. It reinforces the centrality of health as both a determinant and consequence of economic activity, and highlights the necessity of integrative, interdisciplinary approaches to studying these relationships. By quantifying the economic penalties associated with mental illness, obesity, and the effect of relative disadvantage on health, the research offers practical guidance for policymakers seeking to promote healthier, more inclusive, and more economically productive societies.

Chapter 1

The Effect of Mental Health on Early Retirement Decisions: Evidence from Australia

AARON KATZ

Abstract

Health and labour supply are interconnected; However, research has predominantly focused on the impact of physical health, leaving a gap in understanding the role of mental health problems. This study addresses this gap by examining the effect of mental health on early retirement decisions using data from the Household, Income, and Labour Dynamics in Australia (HILDA) Survey. We use both linear probability models and a discrete-time hazard approach. While linear models estimate the average effect, the discrete-time hazard model tracks initially employed individuals aged 50 to 64 over time until they retire early or reach retirement age. To mitigate potential bias arising from the timing of reporting of mental health and retirement decisions, lagged measures of mental health are applied, with respect to the temporal sequence of events. To address measurement bias, the association between our derived mental health variable and other objective psychiatric measures is examined. Furthermore, we include the death of a close friend as an instrument for mental health status, helping us validate and strengthen causal findings of our study. Lastly, we examine whether unobserved heterogeneity poses a problem in our analysis by estimating models with and without unobserved heterogeneity. Our findings indicate a significant and positive causal impact of poor mental health on early retirement decisions, which is also supported by the nonlinear analysis. To explore potential gender heterogeneity, separate analyses are conducted for males and females. The observed differences in the results between the two groups support the assumption of gender-specific effects. These findings suggest that poor mental health has a significant and potentially causal impact on premature exit from the labour market, particularly among men. The results highlight the importance of effective mental health management in supporting longer working lives.

Keywords: Mental health, labour supply, ageing, early retirement, HILDA, discrete-time hazard models, 2SRI

JEL Codes: I12; I15; J14; J26

1.1 Introduction

Mental health issues represent a growing concern for labour markets worldwide, significantly influencing employment outcomes across sectors. Poor mental health affects labour market participation through several channels. Affected individuals may experience increased absenteeism due to illness, medical appointments, or treatment regimens (Bloom and Canning, 2000). Symptoms, such as depression or mood instability, can impede regular work attendance and reduce motivation to invest in human capital (Fadare et al., 2023; Tompa, 2002). The resulting decline in labour productivity and economic engagement may encourage earlier retirement, particularly if individuals anticipate a reduced lifespan and wish to maximise leisure time from accumulated wealth. Additionally, presenteeism, or diminished productivity while at work, may further erode performance, ultimately leading to premature labour market exit (Chatterji et al., 2007). Mental health problems may also increase reliance on non-wage income sources, such as welfare benefits, thereby potentially reducing the financial incentive to remain employed (Disney et al., 2006). However, this relationship is not uniform; some individuals may instead choose to remain in employment longer in order to meet elevated healthcare expenses associated with managing chronic mental health conditions (Bryan et al., 2022; Frijters et al., 2010; Hamilton et al., 1997; Ngui et al., 2010).

In Australia, the economic burden of poor mental health is substantial. Estimates by the Australian Bureau of Statistics suggest that mental health-related work absences cost approximately AU$60 billion annually (Australian Bureau of Statistics, 2008). Furthermore, Lee et al. (2017) estimated that productivity losses stemming from depression, anxiety disorders, and substance use disorders amounted to AU$11.8 billion in 2007, with additional fiscal implications of AU$1.2 billion in lost income tax revenue and AU$12.9 billion in welfare expenditures. For individuals diagnosed with psychosis, Neil et al. (2014) calculated productivity losses at AU$40,941 per person, with further indirect costs, such as those associated with supported employment and non-governmental services, totalling an additional AU$14,642 per person.

This paper makes three key contributions to the literature on mental health and labour supply. First, it tries to provide causal evidence on the effect of mental health on early retirement decisions among older individuals in Australia, addressing an underexplored area in comparison to the well-established literature on physical health and labour market exits. Mental health conditions, often less visible than physical ailments, may receive disparate treatment from employers and policymakers (Ngui et al., 2010). For example, in a similiar economy, the United Kingdom, one in four people reported a mental disorder in

2016 (Alderwick and Dixon, 2019; McManus et al., 2016), yet the National Health Service allocated only 12.5% of its total budget to mental health in 2017/18, with a target of 16.2% by 2022/23 (Baker and Kirk-Wade, 2023).

As Australia's population ages, with those over 50 projected to increase from 12% to 22% between 2015 and 2050 (World Health Organisation, 2022), understanding the causes of early retirement becomes critical for economic planning. An ageing workforce can strain public finances and healthcare systems, increase dependency ratios, and necessitate adjustments to retirement policies (Liddiard, 1978). Determining whether early retirement is driven by health or financial considerations has significant implications for public policy (Frijters et al., 2010). Distinguishing mental from physical health conditions helps in targeting interventions and tailoring incentive schemes that can prolong workforce participation.

To investigate the causal relationship between mental health and early retirement, the empirical analysis proceeds in two stages. First, a linear cross-sectional framework is applied, beginning with Ordinary Least Squares (OLS) estimation and followed by a Two-Stage Least Squares (2SLS) specification. The 2SLS model uses the lagged death of a close friend as an instrument for mental health. Unlike the more commonly used death of a family member, which may directly affect retirement through caregiving or financial implications, this instrument is assumed to influence mental health without directly impacting retirement decisions, helping to address endogeneity in a transparent and interpretable way.

To establish a benchmark and enable meaningful comparisons, we begin by estimating a linear probability model. This serves as a baseline against which the results of subsequent nonlinear analyses can be assessed. Although linear models are intuitive and straightforward, they impose restrictive assumptions and may yield biased estimates. Nonlinear analysis offers an additional framework for modelling time-to-event outcomes like retirement, as it accounts for the dynamic nature of such decisions and handles censoring more appropriately.

The second contribution to the literature, and the second stage of the analysis, extends the investigation using discrete-time hazard models to examine the timing of retirement while explicitly addressing whether unobserved heterogeneity is an issue in our analysis. This approach compares models with and without unobserved heterogeneity, under various distributional assumptions, to assess the robustness of the results. Specifically, a discrete-time hazard model is paired with a two-stage residual inclusion (2SRI) estimator to control for both endogeneity and unobserved confounding (Terza et al., 2008). This more flexible methodology allows for testing whether linear models provide adequate estimates or if failing to account for unobserved heterogeneity and nonlinearity introduces bias, especially when unmeasured factors such as personality traits or job satisfaction may simultaneously influence mental health and retirement decisions.

By leveraging panel data within the discrete-time hazard framework, the longitudinal structure captures individual dynamics over time. While the discrete-time model assumes selection on observables, i.e., that mental health only affects retirement through measured variables, this assumption may be violated if unobserved factors influence both outcomes. To address this, the 2SRI procedure improves causal inference by mitigating bias from such confounding in a panel data setting.

Lastly, the paper explores the role of measurement error and reporting bias in mental health assessment. Unlike many existing studies that rely exclusively on self-reported mental health status, this study compares subjective and objective indicators to account for potential biases in reporting. Social desirability or stigma may lead respondents to underreport psychological conditions, which can result in attenuated estimates of the true effects (Bharadwaj et al., 2017; Brohan and Thornicroft, 2010; Rüsch et al., 2005). By triangulating different data sources and validating mental health measures, the analysis seeks to provide more robust and credible findings.

The results reveal a statistically significant association between poor mental health and an increased likelihood of early retirement. The instrumental variable analysis supports a causal interpretation, with the death of a close friend serving as a valid and relevant instrument, as indicated by the first-stage statistics. When the analysis is extended to a longitudinal framework, the findings remain consistent, with the estimated effect of mental health on early retirement slightly larger than in the non-IV models. This pattern suggests that failing to address endogeneity may lead to an underestimation of the true impact of mental health.

The importance of correcting for endogeneity is further highlighted by the notably larger coefficient, more than six times greater, in the 2SLS model compared to the baseline OLS estimate. Moreover, accounting for unobserved heterogeneity in the longitudinal models does not substantially alter the results, with estimates remaining similar across models with and without frailty.

Importantly, the results indicate gender-specific differences in the mental health and retirement relationship. For males, poor mental health significantly increases the likelihood of early retirement across most specifications, whereas for females, the effect is weaker and not statistically significant in some models. These findings offer important insights for policymakers aiming to extend working lives, improve mental health support, and mitigate the fiscal challenges posed by an ageing population. This study contributes to the existing literature by providing robust evidence from Australia, applying rigorous econometric methods, and underscoring the significance of gender and measurement issues in mental health research.

The outline of the paper is as follows. The next section (*section 1.2*) presents a review of the existing literature that examines the impact of mental health on labour market outcomes and early retirement decisions. *Section 1.3* describes the empirical methods, including linear probability model, two-stage least squares, discrete-time hazard models with and without unobserved heterogeneity, validity checks, and the inclusion of the two-stage residual inclusion approach within a discrete-time hazard model. This is followed by the data description in *section 1.4*. *Section 1.5* presents the results, while *section 1.6* contains the discussion and conclusion.

1.2 Background and Literature Review

Mental health disorders are diverse and vary in their symptoms, treatments, diagnosis, and outcomes (American Psychiatric Association, 2013). The severity of mental illness varies greatly from common illnesses, such as general anxiety, mood swings, eating disorders, to more severe mental disorder e.g., schizophrenia. There is a wide range of literature within different fields of research that show mental health disabilities are linked to poorer labour market outcomes, in addition to affecting social life, such as terminating relationships, loneliness, and a greater likelihood of being involved in a crime (Bartel and Taubman, 1979, 1986).

There is a substantial amount of previous research on the relationship between health and labour market outcomes, such as economic inactivity or retirement. However, most of the research focuses on the effect of physical or general health, without controlling for mental health specifically. Few papers which analyse the effect of mental health and labour market outcomes, such as employment status (Alexandre and French, 2001; Bryan et al., 2022; Chatterji et al., 2011; Ettner et al., 1997; Frijters et al., 2010; Lu et al., 2009; Sainsbury et al., 2008), income (Chatterji et al., 2011; Ettner et al., 1997) and work hours (Ettner et al., 1997). Stern (1989) observed the effect of health on retirement decisions, while using longitudinal data and observing health using subjective measures.

The empirical analysis is grounded in the standard life-cycle labour supply framework, in which individuals allocate time between work, leisure, and health investment to maximise lifetime utility (Grossman, 1972). Within this framework, health, including mental health, affects both the utility derived from work and the productivity of labour. To empirically assess this relationship, we estimate baseline linear models (LPM and 2SLS) where employment status serves as a reduced-form representation of labour supply decisions, and mental health is treated as a potentially endogenous determinant of labour market participation.

Stern (1989) was one of the first to identify a negative and statistically significant effect of health on retirement and labour supply, suggesting that poor health leads to a reduced labour market participation. Similarly, Bryan et al. (2022) found that poor mental health decreases the probability of being employed, while Lu et al. (2009) reported that a decline in average mental health is associated with a significant reduction in both employment rates and annual income. Their findings also indicate that the mental health index has a positive and significant effect on the likelihood of being employed. Despite these contributions, recent studies have focused mainly on general labour market outcomes, with limited attention to older individuals and specific outcomes such as early retirement decisions, exceptions include the work of Zucchelli et al. (2010) and Disney et al. (2006), who specifically examine the effect of general health on labour market outcomes across older workers.

Previous studies of the effect of mental health on labour market outcomes, based mainly on cross-sectional data to examine this relationship (Bazzoli, 1985; Bound et al., 1999; Cai and Kalb, 2007; Siddiqui and Ali Shah, 1997; Zhang et al., 2009). However, the use of cross-sectional data poses limitations in capturing the dynamic effects of health on labour supply and addressing the issue of possible endogeneity through unobserved personal characteristics. To overcome these limitations, longitudinal data has been recommended to mitigate selection bias and gain a deeper understanding of individual behaviour (Chatterji et al., 2011; Nerlove, 2005).

The discrete-time hazard analysis, widely used in the biomedical field and more recently in health economics research, provides valuable insights into the impact of health on workforce outcomes (Bünnings and Tauchmann, 2015; Disney et al., 2006; Zucchelli et al., 2010). This approach offers several advantages. Firstly, it allows greater flexibility in modelling dynamics and exploring variations in the impact of health on labour market transitions according to current employment status (Disney et al., 2006). Secondly, it enables examination of the timing and sequencing of events, offering a more nuanced understanding of the process leading to early retirement. By modelling the transition probabilities over discrete-time intervals, we can capture the dynamic nature of the decision-making process and account for changes in mental health status over time. Finally, the discrete-time hazard approach accommodates time-varying covariates, such as changes in mental health status, marital status, and other relevant factors, allowing for a more comprehensive analysis of the determinants of early retirement decisions.

Building on this methodological foundation, previous empirical studies using longitudinal data have shown that poor self-reported general health is strongly associated with early retirement among older adults (Disney et al., 2006; Zucchelli et al., 2010). For instance, Disney et al. (2006) and Zucchelli et al. (2010) apply discrete-time hazard models using the HILDA dataset, adopting a non-parametric approach to the hazard function. Our study extends this line of research by employing the same discrete-time hazard framework described by Jenkins (1995), but focusing specifically on mental health rather than general health. Furthermore, consistent with Disney et al. (2006); Jenkins (1995, 1997), we account for unobserved heterogeneity to ensure robust inference.

While the discrete-time hazard framework provides a powerful tool to analyse transitions into early retirement, it also brings to the forefront an important econometric concern, namely endogeneity. In studying the relationship between mental health and retirement, endogeneity may arise due to measurement error in self-reported health, the potential for reverse causality between mental health and labour market participation, or omitted variables correlated with both. Addressing this issue is crucial to ensure that estimated effects reflect a causal rather than merely associative relationship.

The issue of endogeneity in our study is particularly relevant due to the potential presence of unobserved characteristics and events associated with both mental health and early retirement decisions (Bryan et al., 2022). To mitigate potential endogeneity bias, our study draws on strategies that exploit exogenous variation in health. One common approach involves identifying health shocks that are plausibly independent of unobserved individual characteristics (Bound et al., 1999; Disney et al., 2006). In this study, we define a health shock by incorporating both lagged and initial-period health. By conditioning on initial health, the coefficient on lagged health can be interpreted as a deviation from an individual's underlying health stock (Disney et al., 2006), thereby helping to control for unobserved, time-invariant health-related heterogeneity (Jones, 2009).

We also control for the latter to minimise omitted variable bias, as both mental and physical health can influence retirement decisions. Specifically, we construct a variable representing a negative physical health shock, defined as the difference between expected and actual self-reported physical functioning scores. Following Apouey et al. (2019), we create a binary indicator that takes the value of one when the observed decline in health is greater than one standard deviation relative to the individual's expected change in health status. This measure helps isolate mental health effects that are not merely reflections of concurrent changes in physical condition.

While exogenous shocks can reduce endogeneity, they may not fully eliminate it. To strengthen causal identification, instrumental variable (IV) techniques are widely employed in health and labour economics. The work by Angrist et al. (1996) establishes a comprehensive framework for identifying causal effects using IVs, outlining the necessary assumptions, such as relevance and the exclusion restriction, and the estimation techniques required for valid inference. This framework has since guided a broad range of empirical applications in health and labour economics, highlighting both the potential and the limitations of IV strategies, including issues of weak instruments, overidentification, and bias arising from unobserved confounding variables (Alexandre and French, 2001; Angrist and Pischke, 2009; Ettner et al., 1997; Frijters et al., 2010; Hamilton et al., 1997; Terza et al., 2008; Wooldridge, 2010; Zhang et al., 2009).

Building on these methodological insights, our study considers instruments proposed in the literature for mental health. One commonly used strategy relies on the death of a family member, under the assumption that such an event affects mental health but is not directly related to labour market outcomes (Böckerman et al., 2022; Burrell et al., 2022). However, applying this strategy in the context of employment decisions introduces additional complexity. The exclusion restriction may be violated if the death of a family member influences labour supply through alternative pathways, such as inheritance or changes in household income. For instance, the death of a spouse may compel the individual to continue working for financial reasons, thereby confounding the relationship between mental health and employment.

To address these challenges, we adopt a two-step strategy that strengthens identification. First, we refine our measurement of mental health by regressing alternative indicators of psychological well-being on the main mental health variable as a robustness check, helping assess potential measurement bias. Second, we apply an IV approach that uses the death of a close friend as an instrument for mental health, a shock that is plausibly exogenous to labour market outcomes. The discrete-time hazard model is then estimated under alternative distributional assumptions, both with and without unobserved heterogeneity, and incorporates a two-stage residual inclusion (2SRI) procedure to account for endogeneity within the nonlinear framework (Bound et al., 2001; Butler et al., 1987; Disney et al., 2006; Kessler et al., 2002; Zucchelli et al., 2010).

The 2SRI method provides an alternative to the conventional IV approach when applied to nonlinear models. In the context of discrete-time hazard analysis, the 2SRI approach has been applied to estimate the effect of a time-varying exposure or treatment on a time-to-event outcome (Basu and Rathouz, 2005; Garrido et al., 2012). Previous studies have used the

2SRI approach to estimate the effect of various health-related interventions, such as diabetes treatments, drug treatments, and coverage of health services, on hazard outcomes (Mery et al., 2016; Tchetgen et al., 2015; Ying et al., 2019). The 2SRI approach has the advantage of accounting for the endogeneity of the treatment variable, which can improve the accuracy of the estimated treatment effect, but it may still remain biased.

The use of IVs in discrete-time hazard analysis can be challenging due to the need for large sample sizes and the potential for weak instruments. Weak instruments can result in biased estimates and large standard errors, leading to incorrect inference (Angrist et al., 1996; Terza et al., 2008). In addition, the use of IVs assumes that the exclusion restriction holds, meaning that the IV only affects the outcome through its effect on the endogenous variable, and not through any other pathway. This assumption can be difficult to verify in practice. Studies have highlighted the effectiveness of the IV approach in addressing endogeneity and providing insights into the effects of variables on survival outcomes (Ettner et al., 1997; Tchetgen et al., 2015; Terza et al., 2008).

Our study builds on the econometric framework established by Terza et al. (2008) and Terza (2018), which formalises the use of the 2SRI estimator in nonlinear models. In this context, a valid instrument must satisfy two key criteria: (1) it must be strongly correlated with the endogenous explanatory variable (relevance), and (2) it must be conditionally independent of the outcome, given the endogenous regressor and other covariates (exogeneity). Although, these conditions are similar to those in 2SLS, 2SRI differs in how endogeneity is addressed in nonlinear models: by including the first-stage residual as an additional regressor in the second-stage model, rather than using predicted values as in 2SLS.

Literature on the effect of mental health on labour market outcomes in Australia is infrequent, and often limited to a sub-sample analysis of males. Cai and Kalb (2007) use HILDA survey to examine the effect of health and labour participation. Using a simultaneous equation model for working-age participants to control for potential endogeneity of health, their findings show that health has a positive effect on retirement. Wilkins (2004) and Brazenor (2002) use cross-sectional data to analyse the effect of 10 different disabilities on labour market outcomes of the older population, such as income level and employment status. Wilkins (2004) shows that, on average, disability reduces the likelihood of labour market participation with different effects for males and females. However, Brazenor (2002) finds that the impact varies depending on the disability type.

Bubonya et al. (2019) focus on depressive symptoms as a proxy for having a mental health disorder and its effect on employment status, using the first 14 waves in HILDA. To measure depressive symptoms, they use the same mental health variable available in HILDA as in our paper. As their primary focus is on anxiety and mood disorders, they assign different weights to the five items used to construct the derived mental health variable. Based on this approach, scores below 60 are coded as moderate symptoms, while scores below 52 are considered indicative of severe symptoms.

1.3 Methodology

Existing empirical research has adopted two main strategies to address the structural endogeneity inherent in analysing the relationship between health and labour market outcomes. One common approach involves identifying the model through theoretically motivated exclusion restrictions, this requires finding instruments that affect health but not employment. However, the validity of such instruments is often difficult to defend in practice. An alternative strategy relies on exploiting the temporal ordering of events to mitigate reverse causality, typically by modelling the effects of lagged health and employment outcomes (e.g., Olesen et al., 2013; Steele et al., 2013).

In line with this, recent work by Bubonya et al. (2019) using Australian panel data (HILDA) examines the bidirectional relationship between depressive symptoms and employment. They apply linear fixed-effects and dynamic panel models to control for unobserved heterogeneity and to explore feedback mechanisms over time. Their analysis confirms a robust negative effect of poor mental health on employment probabilities, while also documenting that job loss exacerbates psychological distress, reinforcing concerns about simultaneity bias. These insights inform our empirical strategy. We begin by estimating a linear probability model to establish a baseline association between mental health and early retirement. We then proceed to an instrumental variable framework to address potential endogeneity arising from omitted variables and reverse causality. Finally, to capture the timing of retirement decisions in relation to health, we include a discrete-time hazard model.

1.3.1 Baseline Linear Model: LPM and 2SLS

To gain an initial understanding of the relationship between mental health and early retirement decisions, we first estimate a linear probability model (LPM), in which an outcome variable of reported retirement is regressed on lagged mental health and a set of control variables using a pooled regression framework for individuals aged 50–64. This baseline model offers a straightforward interpretation of the effects. However, this approach does not account for the timing of retirement events, the longitudinal structure of the data, or the possibility of censoring over time. To more appropriately address these dynamic aspects, the analysis is extended using discrete-time hazard models, which explicitly model the retirement decision as a time-to-event process.

The baseline LPM model captures the relationship between mental health and the probability of retirement. The specification is given by:

$$Y_{it} = \alpha_0 + \alpha_1 m_{it-1} + \alpha_2 X_{it} + \varepsilon_{it} \tag{1.1}$$

where Y_{it} is a binary indicator equal to 1 if individual i reports retirement at time t and 0 otherwise, m_{it-1} denotes the lagged mental health score with lower values indicating poorer mental health and higher values reflecting better mental health, and X_{it} is a vector of socio-demographic controls. The error term ε_{it} captures unobserved factors affecting retirement. Time subscripts are included; however, we estimate the model by pooling observations across individuals and time, effectively treating the data as if it consists of $i \times t$ independent units.

Although LPM offers a straightforward interpretation, the coefficient α_1 measures the change in the probability of retirement associated with a one-unit change in lagged mental health score, it has several limitations. Notably, the LPM can produce predicted probabilities outside the $[0,1]$ interval and assumes linearity in probabilities, which might be restrictive given the binary outcome. More importantly, although OLS provides a straightforward estimate, it is likely to suffer from endogeneity due to reverse causality or omitted variable bias, where poor mental health causes and is caused by retirement decisions. For example, retirement could itself impact mental health, creating reverse causality bias.

To mitigate endogeneity concerns, we implement a Two-Stage Least Squares (2SLS) instrumental variables strategy. We use the death of a close friend in the previous period to mental health (t-2) as an instrument for lagged mental health (t-1). This instrument is relevant because bereavement is empirically linked to worsened mental health outcomes (Frijters et al., 2010). Under the key assumption that the death of a close friend affects retirement decisions only through its effect on mental health (exclusion restriction), the 2SLS estimates provide a more reliable causal effect of mental health on retirement probability. The first-stage predicts mental health using the instrument, and the second-stage regresses retirement on the predicted mental health and covariates, thus purging the mental health variable of endogeneity bias.

The first-stage of the 2SLS model is specified as:

$$m_{it} = \pi_0 + \pi_1 Z_{it-1} + \pi_2 X_{it} + \eta_{it} \tag{1.2}$$

where Z_{it-1} is an indicator for the death of a close friend reported at time $t-1$, referring to an event that occurred between $t-2$ and $t-1$. This variable therefore instruments $m_{i,t-1}$ for the outcome observed at time t, capturing the delayed effect of bereavement on mental health. The additional lag allows for the time it can take for grief to translate into measurable mental health deterioration.

The second-stage uses the lagged predicted mental health score from the first stage, $\hat{m}_{it-1}$, to estimate its causal impact on early retirement:

$$Y_{it} = \beta_0 + \beta_1 \hat{m}_{it-1} + \beta_2 X_{it} + v_{it} \tag{1.3}$$

This IV strategy allows us to recover a local average treatment effect (LATE), interpreted as the causal effect of mental health on retirement among individuals whose mental health was affected by the bereavement shock. A statistically significant and negative β_1 would indicate that better mental health causally decreases the likelihood of retirement.

To assess instrument strength, we report the first-stage F-statistic. A value above the conventional threshold of 10 suggests a strong instrument. Additionally, validity checks are conducted using alternative specifications of bereavement timing and sub-group analyses by gender.

To account for unobserved individual heterogeneity that may bias the estimates in the baseline linear model, we also estimate a fixed-effects panel regression. This approach controls for time-invariant individual characteristics (e.g., personality traits, early-life conditions) that may jointly influence both mental health and retirement decisions. The fixed-effects specification is given by:

$$Y_{it} = \alpha_1 m_{it-1} + \alpha_2 X_{it} + \mu_i + \varepsilon_{it} \tag{1.4}$$

where μ_i captures unobserved individual-specific effects, and ε_i denotes the idiosyncratic error term after accounting for μ_i . By differencing out time-invariant unobservables, the fixed-effects estimator helps to address bias from omitted variables that are constant over time. However, it does not account for time-varying endogeneity, which is addressed in instrumental variable strategies.

We include the linear approach alongside the discrete-time hazard analysis to establish a baseline and enable meaningful comparison. Although linear probability models offer a straightforward starting point for examining the relationship between mental health and early retirement, they come with notable limitations. Chief among these are the assumptions of constant marginal effects and the risk that predicted probabilities may fall outside the valid range of 0 to 1. These constraints become particularly problematic when modelling binary or time-to-event outcomes such as retirement.

Furthermore, linear models are limited in their ability to analyse retirement behaviour, as they do not track individuals over time and therefore fail to fully exploit the longitudinal structure of the HILDA dataset. These models also struggle to capture the timing of retirement decisions and the presence of time-related features such as censoring. To address these limitations, we include a discrete-time hazard modelling framework that explicitly follows individuals over time and accounts for the dynamic nature of retirement behaviour. This nonlinear approach provides a more flexible and accurate representation of the timing of retirement, while appropriately handling duration dependence and potential right-censoring in the data.

1.3.2 Discrete-time Hazard Model

We apply the discrete-time hazard model and assume that individuals become at risk of early retirement at the age of 50 years old, which is in line with the previous literature (Butterworth et al., 2006; Melzer et al., 2004). Setting the pattern of duration dependence is essential prior to analysing the model, and thus we include an additional 12 age dummies to capture the duration dependence by the age of the individuals, starting from 53 to 64 years old. In our sample of HILDA data, which is based on the number of retirees within each age group, the actual risk of retiring early started when participants are 53 years and older, since the number of reported retirements between the ages 50 and 52 is minimal in our sample.

Our research focuses on understanding how mental health influences the decision to retire. In order to observe individuals who are at risk of retirement, we define a stock sample based on the definition of Jenkins (1995). Over the course of the study, only 45 individuals were present for all 15 waves due to sample attrition, retirement, and death. Our retirement models are estimated on complete sequences of observations, using information up to the wave of first exit if an individual leaves the panel and returns in later waves.

We create additional rows of observation per individual based on the time that this individual was at risk of reporting early retirement in our data. This requires that in general, each respondent, i, contributes T_i rows, where T is the number of time periods i was observed at risk of failure. Jenkins (1995) method enables all periods prior to selection to be ignored. This approach relies on only focusing on the relevant time restricted data for observing the effect of interest. Additionally, the data needs to be rearranged and conditioning on stock sampling prior estimation and need to have an unbalanced data format.

We follow the same notation as Jenkins (1995), so time, t, equals to the initial period, τ, $(t = \tau)$ when the individual enters the analysis. Each respondent i contributes s_i years of risks while they are still in employment in the interval between the initial period τ and s_i, so the time when retirement occurs is $(t = \tau + s_i)$. Here, we only rely on the age of an individual reported in the first wave in the survey for evaluating the retirement age. Additionally, we follow the age variables on wave 2 onward. The age and our derived time variable may differ as the timing of the surveys varies between the 18 waves and can cause respondents to report the same age in two consecutive waves, if the timing of the later wave was before their birthday. We assume throughout those individuals aged one year between waves of data.

At the end of the sample, each individual will either have early retired (where $\delta_i = 1$) or will be censored and thus will still be working during the last wave in our sample, $(\delta_i = 0)$. The age of the individual $t_i = \tau_i + s_i$ will be the retirement age in a complete duration data or the final age of observation if $\delta_i = 0$. As we ignore all the periods prior to selection into the stock sample (Jenkins, 1995), we initially selected only those who are participating in the labour market at their first wave, and therefore drop all individuals who reported their occupational status as retired in the wave when joining the survey.

The dependent variable is a binary indicator for whether the participant remained in or left the labour market through retirement in that wave. We do not allow individuals to report multiple retirements by coming back to the labour market and then leaving again. That is retirement is considered as an absorbing (permanent) state, where individuals cannot return to the labour market.

Consider the discrete-time hazard rate, for individual i, with the conditioned probability of retirement at age t, given by:

$$h_{it} = P[T_i = t | T_i > t; X_{it}, m_{it-1}] \tag{1.5}$$

Where T_i is a discrete random variable, which represents the last derived age observed in the wave at the end of the sample time period $(t = \tau + s_i)$. X_{it} is a vector of socio-demographics covariates based on previous literature, and including marital status, gender, number of individuals in the household, education level, unemployment rate in the local region, and household income, and m_{it-1} is mental health score of individual i at time t-1.

The analysis is conditioned a stock sample, implying that all periods prior to the selection period can be ignored, given that individuals are observed to be in employment at the beginning of the sample time period. The conditional probability of observing the event history of an individual without a complete sequence of responses for the whole sample period is:

$$P(T_i > t | T_i > \tau - 1) = \prod_{t=\tau}^{\tau+s_i} (1 - h_{it}) \tag{1.6}$$

The conditional probability of observing an individual with a complete sequence between τ and the time of the interview is:

$$P(T_i = t | T_i > \tau - 1) = h_{i\tau+s_i} \prod_{t=\tau}^{\tau+s_i-1} (1 - h_{it}) \tag{1.7}$$

Equation(1.7) can be simplified by:

$$(\frac{h_{i\tau+s_i}}{1 - h_{i\tau+s_i}}) \prod_{t=\tau}^{\tau+s_i} (1 - h_{it}) \tag{1.8}$$

Combining Equation(1.7) and Equation(1.8), for the corresponding log-likelihood (with and without a complete sequence) for the whole sample yields:

$$logL = \sum_{i=1}^{n} \delta_i log(\frac{h_{i\tau+s_i}}{1 - h_{i\tau+s_i}}) + \sum_{i=1}^{n} \sum_{t=\tau}^{\tau+s_i} log(1 - h_{it}) \tag{1.9}$$

The log likelihood function in Equation(1.9) depends on the labour market status of the individual i at the end of the sample time period. The individual can retire before the sample time period, $\delta_i = 1$, or to have a complete duration spell if $\delta_i = 0$.

For individuals who stay in the labour market until the last observed wave, denoted by $y_{it} = 0$, the value remains the same for all spell periods. On the other hand, for individuals who exit, denoted by $y_{it} = 0$ for all periods except the exit period, where it becomes $y_{it} = 1$.

The log-likelihood can be simplified by replacing δ_i with y_{it} in Equation 1.9:

$$logL = \sum_{i=1}^{n} \sum_{t=\tau}^{\tau+s_i} y_{it}\, log(\frac{h_{it}}{(1 - h_{it})}) + \sum_{i=1}^{n} \sum_{t=\tau}^{\tau+s_i} log(1 - h_{it}) \tag{1.10}$$

The hazard rate is the rate of retirement at time t and measures how probable an observation is to retire as a function of age of the individual conditioned on surviving to t-1. The hazard rate is defined by applying a complementary log-log hazard rate, as follows:

$$h_{it} = 1 - exp(-exp(\beta_1 m_{it-1} + \beta_2 X_{it} + \theta_t)) \tag{1.11}$$

Where θ_t is the discrete-time baseline hazard in our model. We include age dummies for every year at risk of retiring early, starting from 53 to 64 years old, as the actual risk of retiring started at 53 years old. The age dummies between 50 and 52 are not included as they are captured in the model's constant. Therefore, we use the semi-parametric form of our hazard model, where $h_0(t)$ specified as a step function, using dummy variables for each year of age. We also estimate sub-samples of males and females separately as it is likely that the response of labour supply to health shocks differs by gender (Stock and Wise, 1988).

1.3.3 Unobserved heterogeneity

Unobserved heterogeneity refers to individual-specific characteristics that are not captured by the observed covariates but still influence the outcome variable, such as differences in preferences, motivation, or unmeasured health conditions (Lemeshow et al., 2011). Ignoring these latent factors in a hazard model can lead to biased estimates of the baseline hazard and the coefficients of interest, as the model may attribute part of the variation in hazard time to observed covariates rather than to unobserved differences across individuals. This can result in premature censoring and distortions in the estimated duration dependence, producing either an overestimation or underestimation of the true effect of covariates (Balan and Putter, 2019; Heckman and Singer, 1984; Lancaster, 1990; Nicoletti and Rondinelli, 2010).

In general, neglecting unobserved heterogeneity or misspecifying its distribution can bias the estimation of the hazard rate and the coefficients associated with explanatory variables (Heckman and Singer, 1984). Individuals with higher unobserved risk may exit the labour market earlier, leading to spurious negative duration dependence (Lancaster, 1990). Moreover, when the form of heterogeneity is incorrectly modelled, the estimated relationship between regressors and the hazard rate may not reflect the true behavioural mechanisms (Nicoletti and Rondinelli, 2010).

These biases are generally more problematic in duration models than in linear regressions, where the consequences of unobserved heterogeneity are less severe if it is uncorrelated with the regressors. In hazard settings, however, the nonlinearity of the model amplifies the effects of omitted heterogeneity, threatening the validity of the estimates (Jenkins, 1995; Lancaster, 1990; Nicoletti and Rondinelli, 2010).

To assess the sensitivity of our estimates, we also follow Nicoletti and Rondinelli (2010) who use Monte Carlo simulations to evaluate the impact of misspecifying unobserved heterogeneity. Their findings suggest that discrete-time hazard models, particularly those using the complementary log-log specification, are relatively robust to certain forms of misspecification, though accounting for frailty can still improve accuracy.

In our empirical analysis, we begin by estimating a complementary log-log model without frailty, followed by specifications that incorporate unobserved heterogeneity using a gamma-mixture distribution (Jenkins, 1995, 1997). Model fit is compared using the Akaike Information Criterion (AIC) to evaluate whether including unobserved heterogeneity meaningfully alters the results. This analysis directly informs the following section, where we extend the framework to include a two-stage residual inclusion (2SRI) approach to address potential endogeneity.

1.3.4 Two-stage Residual Inclusion in Discrete-time Hazard Settings

The methodology used in this study is a 2SRI approach in a discrete-time hazard analysis context. The 2SRI estimator is an alternative to the two-stage least squares (2SLS) estimator but applicable to non-linear settings, while in linear models 2SRI is equivalent to 2SLS. The 2SRI incorporates the endogenous variable in the second-stage in addition to the predicted residual term from the first-stage to further account for endogeneity. The 2SRI estimator relies on previous literature, including the control function approach proposed by Heckman and Robb Jr (1985) and the residual inclusion method introduced by Hausman and Taylor (1981) and Terza et al. (2008). The 2SRI has been shown to have desirable properties in a range of empirical contexts (Angrist et al., 1996).

In the first-stage, the model is estimated with the instrumental variable as the independent variable and the endogenous variable as the dependent variable. The predicted values of the residuals from this model are then used in the second-stage as an additional covariate to estimate the effect of the exposure variable on the hazard outcome. The 2SRI approach can improve the precision of the estimate and reduce bias due to unobserved confounding, but its validity relies on "strong" instrumental variables and the assumption that the residual from the first-stage is uncorrelated with the unobserved confounders.

We consider a discrete-time hazard model specifications mentioned in Equation(1.10). Hence, the first-stage is further specified as:

$$m_{it} = \alpha_0 + \alpha_1 Z_{it-1} + \alpha_2 X_{it} + u_{it} \tag{1.12}$$

where u_{it} is a mean zero residual error independent of Z, given X_{it}.

In the first-stage, we regress the endogenous covariates on the exogenous instrumental variable Z, and estimate it using an ordinary least squares model (OLS). We obtain the predicted residuals, $\hat{u}_{it}$, of the endogenous covariates from Equation(1.12).

We assume that the instrumental variable, death of a close friend (Z), used in Equation(1.12) is exogenous, meaning it is uncorrelated with the error term in the outcome of the discrete-time hazard model in Equation(1.12). Additionally, given that the effects of a close friend's death on mental health may not be immediately observable and may take some time to materialise, we account for this potential time lag by including a lagged one-period death of a close friend as an instrumental variable in our analysis.

Moreover, we assume that the instrumental variable used in our approach satisfies the exclusion restriction assumption. This indicates that the instrumental variable (lag of death of a close friend), Z, used in the first-stage regression in Equation(1.12) is only affecting the outcome variable (early retirement) through its effect on the endogenous variable, m (mental health score at time t-1), and not through any other unobserved factors. If this assumption is violated, the instrumental variable approach may not be valid and the estimated effect of the endogenous variable, m_{it-1}, on the outcome variable, Y_{it}, may be biased.

We then predict the residuals from the first-stage as:

$$\hat{u}_{it-1} = m_{it-1} - (\hat{\alpha}_0 + \hat{\alpha}_1 Z_{it-2} + \hat{\alpha}_2 X_{it-1}) \tag{1.13}$$

These residuals, $\hat{u}_{it-1}$, represent the unexplained variation in the endogenous variable, m_{it-1} after controlling for the exogenous variation provided by the instrumental variable of the death of a close friend at the previous period (Z_{it-2}).

In the second-stage, we estimate the effect of the predicted residuals $\hat{u}_{it-1}$ on the likelihood of early retirement, Y_{it}, using a cloglog hazard model. After following Aalen (1989) and Wan et al. (2015) where the covariates are allowed to vary over-time, we can show that:

$$logit(P(Y_{it} = 1|\hat{u}_{it-1}, X_{it}, m_{it-1})) = \beta_0 + \beta_1 m_{it-1} + +\beta_2 X_{it} + \beta_3 \hat{u}_{it-1} + v_{it} \tag{1.14}$$

where $P(Y_{it} = 1|\hat{u}_{it-1}, X_{it})$ is the probability of $(Y_{it}) = 1$ given the predicted residuals $\hat{u}_{it-1}$ and the vector of socio-demographic covariates, X_{it}. The parameter β_2 represents the effect of the predicted residuals on $logit(P(Y_{it} = 1))$, controlling for X_{it} and m_{it-1}. v_{it} is a random error term and represents the unobserved factors that affect the outcome variable, Y_{it}. It captures all other factors that might influence, Y_{it}, but are not included in the model, and it is assumed to be uncorrelated with the instrumental variable, Z_{it-2}, other control variables, X_{it}, and m_{it-1} conditioned on $\hat{u}_{it-1}$.

The discrete-time hazard rate is modelled using a complementary log-log specification, such that:

$$h_{it} = 1 - exp(-exp(\beta_1 m_{it-1} + \beta_2 X_{it} + \beta_3 \hat{u}_{it-1} + \theta_{it})) \tag{1.15}$$

where β_1 is the coefficient for mental health status for individual i at time t-1, $\hat{u}_{it-1}$ is the predicated residuals from the first-stage for individual i at time t-1, and other socio-demographic factors for individual i at time t, respectively. θ_{it} is the discrete-time baseline hazard.

We apply the cloglog hazard model that includes a set of exogenous covariates and a discrete-time baseline hazard, θ_{it}, that captures the unobserved factors that affect the hazard of retirement over time. Additionally, we use a semi-parametric approach to model the baseline hazard as a step function, which allows us to estimate the hazard rate at each discrete time t using dummy variables for age to retirement (refer to the next section (section 1.4) for additional explanation).

1.4 Data

This paper uses 18 waves (2001-2018) of the Household, Income and Labour Dynamics in Australia (HILDA[1]) survey. HILDA is a household-based panel study which includes information about socioeconomic characteristics, family dynamics and labour market outcomes. The dataset consists of variables related to individuals and household characteristics, including labour market status, wages, and the health status of individuals. Individuals aged 15 and older are eligible for interview. Both personal and self-assessment questionnaires are used in order to obtain information about the participants.

Participants beyond the state retirement age are excluded, as their retirement decisions may not solely be influenced by adverse mental health conditions. During the time of our sample, the Australian government implemented a policy aligning the female state pension age with that of males, resulting in a notable increase in the female retirement age. Our process involves verifying the age of each female participant in relation to the specific wave, excluding those whose age is equal to or greater than the state pension age during that period.

Individuals with missing employment status (i.e., those who refused to respond or did not provide a response) and those with missing mental health scores, the main variable of interest, are excluded from the analysis. In total, 122,808 observations did not report or refused to report their general health status. Furthermore, 63,956 observations are excluded because individuals were above the statutory retirement age for both genders. Finally, 254,468 observations are excluded for being under the age of 50, leaving a total sample of 44,091 observations. When incorporating lagged mental health, we further exclude individuals who

[1] See https://www.melbourneinstitute.com/hilda/

reported mental health in only one wave, resulting in a final analytical sample of 32,316 observations.

Fig. 1.1 Sample construction and exclusion criteria.

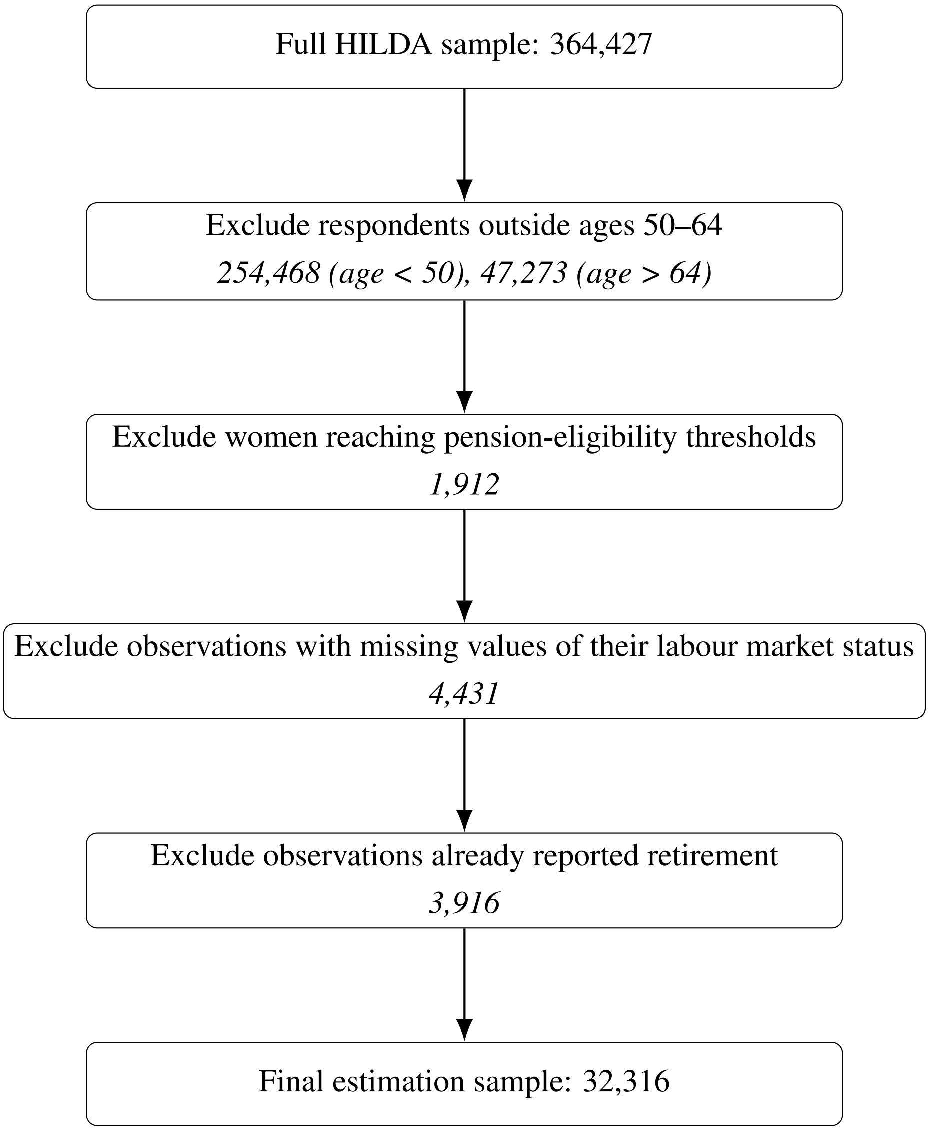

The dependent variable in this study is the mental health component score of the Short-Form health survey (SF-36), as provided in the HILDA dataset. The score is a continuous measure ranging from 0 to 100, with higher scores indicating better mental health status. It is calculated according to the method outlined by (Ware et al., 1993), which weights and combines the individual SF-36 items to produce a standardised score. In this analysis, the score was used as provided in the HILDA dataset, so it was not derived manually. *Figure 1.2* displays the distribution of the mental health variable.

The SF-36 is a self-reported, multidimensional, generic measure of health. Respondents answer 36 short questions regarding their general health status, both physical and mental, relative to other individuals of the same age group, and grade their responses on discrete scales between 1 and 5. The 36 survey items are used to produce an eight-scale profile of functional health and wellbeing as well as a psychometric score. This self-reported questionnaire is based on physical and mental health summary measures, as well as a preference-based health utility index (Ware et al., 1993).

The mental health variable derived from the SF-36 consists of five questions. The items invite the respondents to grade their mental health by replying on a discrete scale to how much they agree with the following statements: been a nervous person, felt so down in the dumps nothing could cheer you up, felt down, been a happy person, felt calm and peaceful. The last two items are reverse-coded so that all items correspond to the same qualitative effect. *Figure 1.2* displays the distribution of the derived mental health variable. The distribution shows that mental health scores are heavily concentrated in the higher range, with most participants scoring between roughly 70 and 90. The distribution peaks around the mid-80s, indicating that a large share of respondents report relatively good mental health. The mean score of 74.92 on a 0–100 scale (where 100 represents the best mental health) reflects this overall skew toward higher values.

Fig. 1.2 Distribution of the derived mental health variable in HILDA from the SF-36

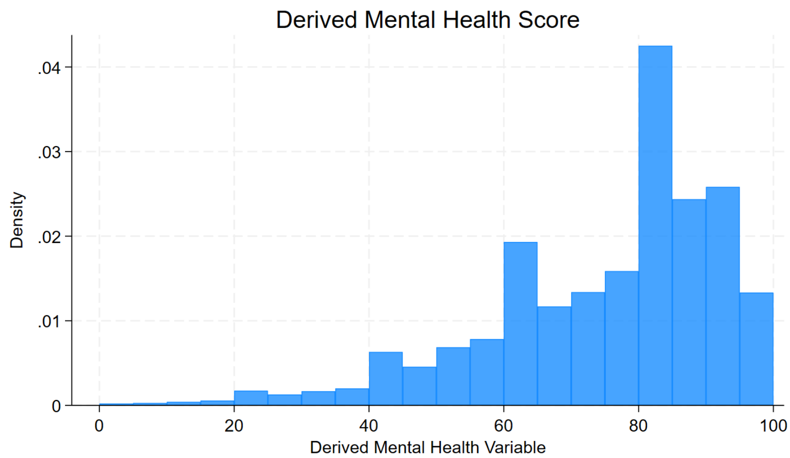

Figure 1.3 presents the kernel density estimates of the derived mental health scores, stratified by retirement status. The graph reveals a clear distinction in the distribution of mental health between retired and non-retired individuals. Specifically, the density of retired individuals is notably higher at lower mental health scores, indicating poorer mental health

within this group. Conversely, the density for non-retired individuals dominates the right tail of the distribution, corresponding to higher (better) mental health scores. This suggests that, on average, retirees report worse general mental health compared to those who remain in the labour force.

Fig. 1.3 Density of Mental Health, by Retirement Status

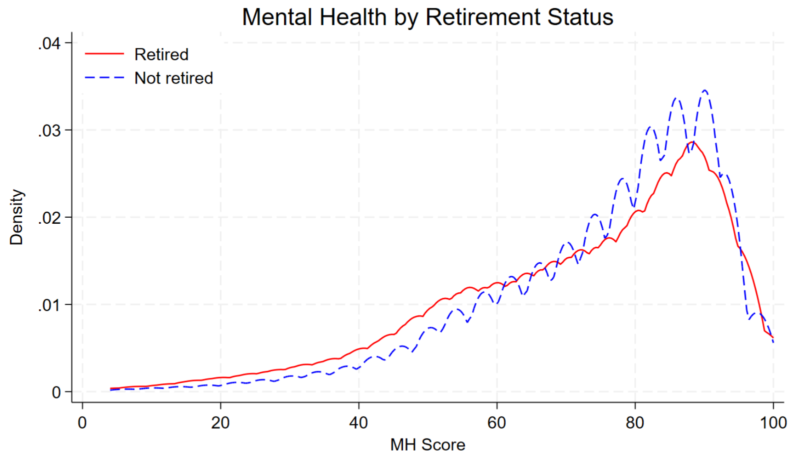

Measurement errors in self-reported mental health may occur due to stigma and potential discrimination (Bharadwaj et al., 2017; Brohan and Thornicroft, 2010; Rüsch et al., 2005). While survey-based self-reported questionnaires are commonly used to measure mental health, the potential for inaccuracy due to social desirability bias raises important considerations (Melzer et al., 2004). Furthermore, when individuals reported mental health, we expect the information to be influenced by evolving social norms and stigma surrounding mental health disabilities (Bharadwaj et al., 2017). However, as societal attitudes towards mental health have evolved in recent years, this shift is reflected in the increased per capita use of public mental health services in Australia (*Figure 1.4*).

Despite the complex challenge posed by measurement error, we mitigate this concern by analysing the correlation between the Kessler Psychological Distress Scale-10 (K-10) questionnaire (Kessler et al., 2002) and our derived mental health variable. The Kessler Psychological Distress Scale includes a 10-item questionnaire, designed specifically for measuring mental health distress caused by anxiety and depressive symptoms following the method of Kessler et al. (2002), where higher scores indicate a greater likelihood of experiencing psychological distress. We run a validity check of the self-derived mental health status variable in HILDA, which is available in all 18 waves. However, the measure of Kessler Psychological Distress Scale is not available in every wave and hence cannot

Fig. 1.4 Number of Medicare Subsidised Mental Health Services Consumed by Australians of all providers per year per 1,000 population (Australian Institute of Health and Welfare, 2021).

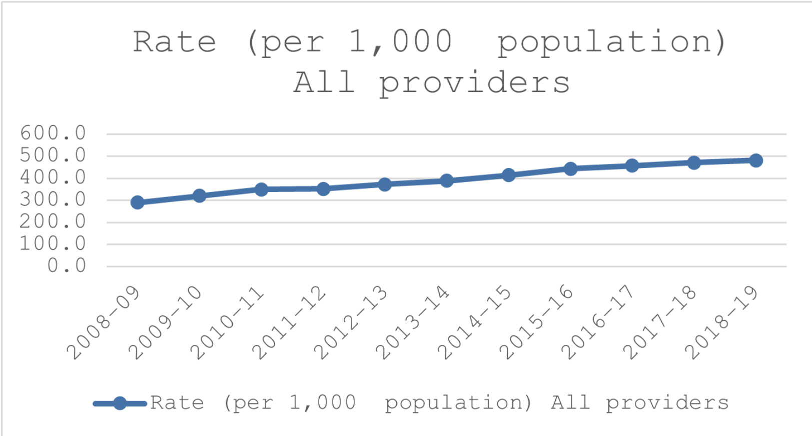

be used as the main measurement of mental health in our analysis. The regression method can provide supporting evidence of using the derived mental health variable available in HILDA, while avoiding a potential measurement error by including a more objective in its measurement.

Applying the K-10 as a benchmark, after recoding so that higher values indicate better mental health, we assess the validity of our self-reported mental health variable in the HILDA dataset, thereby enhancing the robustness of our analysis. Our examination reveals a highly statistically significant relationship between the K-10 scores and our derived mental health variable, exhibiting a substantial correlation between the two measures (*Table 1.1*). The correlation between these two measure reported to be high with more than 80%. This comprehensive analysis helps alleviate potential concerns of measurement error and underscores the reliability of our findings. By including this approach, we align our investigation with the evolving landscape of attitudes towards mental health and enhance the precision of the derived mental health variable included in our analysis.

Table 1.1 OLS regression of the derived variable of mental health on K10

VARIABLES	Derived Mental Health
Kessler	2.203***
	(0.014)
Constant	108.44***
	(0.234)
Observations	13,588
R-squared	0.641
Correlation	-0.801

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

In our paper, an additional regression analysis is applied to examine the association between the derived mental health variable and a binary variable indicating whether individuals have received a formal diagnosis of depression or anxiety by a healthcare professional. This is a formal diagnosis and is more objective and less prone to measurement error. It is important to note that this binary variable is only available in wave 9 and wave 13 and has been queried for a relatively small subset of participants, totalling N=5,436. The regression results, as presented in *Table 1.2*, reveal a statistically significant correlation between the binary variable representing a positive diagnosis of either depression or anxiety and the derived mental health score.

Table 1.2 OLS regression of the derived variable of mental health on depression or anxiety

VARIABLES	Derived Mental Health
Dependanx	-9.470***
	(0.326)
Constant	71.96***
	(0.244)
Observations	5,436
R-squared	0.134
Correlation	-0.365

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

We include mental health lagged one period to exploit the timing of events between shock to mental health on retirement decisions. This removes bias that would occur if mental health status was measured following the retirement decision during any particular wave. By conditioning on a previous reported mental health status, we can be assured that any health "shock" occurred prior to a retirement decision. Lagged health may also be more informative about the decision to retire than contemporaneous health status, simply because transitions take time. It also may take time for an individual to understand that they suffer from a mental health disorder and to adjust to their new health condition, while observing its effect on labour productivity.

We follow Disney et al. (2006); Jones (2009); Zucchelli et al. (2010) and include additional covariates to control for other socioeconomics characteristics that might affect early retirement decisions, such as gender, highest educational attainment, and local unemployment rate. *Table 1.3* shows definitions of all the variables included in our analysis. Additionally, we rescaled the household income by the consumer price index in Australia based on 2001 as the base year, in order to capture inflation and ensure values are expressed in constant Australian dollars.

An additional measure is incorporated by conditioning mental health on an exogenous indicator of a negative physical health shock. Following Apouey et al. (2019), we construct this measure by comparing each individual's expected and actual physical health status across survey waves. Expectations of next-period health are derived from a question in the SF-36

Table 1.3 Variable Definitions

Variable Name	**Variable Label**	**Definition and Measurement**
Retired	Reported to be retired	Binary: 1 = retired, 0 = otherwise
Lagged mental health	Lagged self-reported mental health	Continuous scale (0–36), higher = worse mental health
Derived mental health	Mental health status derived from the self-reported GHQ	Continuous GHQ-derived score (0–36), higher = better mental health
Local unemployment rate	Unemployment rate in major statistical region	Continuous percentage (%)
Household size	Number of in-scope persons in household	Count variable
Degree	Holds an academic degree	Binary: 1 = degree, 0 = otherwise
Married	Legally married	Binary: 1 = married, 0 = otherwise
Initial mental health	Initial mental health status	Continuous scale (0–36), baseline measure
Income	Annual income in Australian dollars ($K)	Continuous, measured in thousands of AUD
Death of a friend	Death of a close friend in the last 12 months (instrument)	Binary: 1 = yes, 0 = no
Age	Respondent's age	Continuous, in years
Negative health shock	Negative health shock of physical functioning	Binary: 1 = experienced shock, 0 = no shock
Kessler	Kessler Psychological Distress Scale-10	Continuous scale (10–50), higher = better mental health (rescaled)
Depandanx	Diagnosis of depression and/or anxiety	Binary: 1 = diagnosed, 0 = not diagnosed
D	Early retirement indicator	Binary: 1 = early retired, 0 = censored
T	Analysis time when record ends	Continuous, measured in years
Wave	Survey wave	Integer, indicates survey round

DV: Dummy Variable

survey contained in the HILDA dataset, where respondents indicate their level of agreement (from 1 to 5) with the statement, "I expect my health to get worse." These responses are coded so that higher values reflect a stronger expectation of health deterioration.

Let PH_{it} denote the self-reported physical functioning score for individual i at time t, and E_{it} represent the expected change in health status for the following wave, derived from responses to the statement "I expect my health to get worse." We first compute the change in actual physical health between two consecutive waves as $\Delta PH_{it} = PH_{it} - PH_{i,t-1}$. We then classify this change into three categories:

- -1 if $\Delta PH_{it} < -1$ standard deviation (decline in health),
- 0 if $|\Delta PH_{it}| \leq 1$ standard deviation (no significant change),
- $+1$ if $\Delta PH_{it} > 1$ standard deviation (improvement in health).

$$S_{it} = \begin{cases} 1 & \text{if the individual expected their health to stay the same or improve } (E_{it} \geq 0) \\ & \text{but experienced an actual decline } (\Delta PH_{it} < -1), \\ 0 & \text{otherwise.} \end{cases}$$

In this definition, $S_{it} = 1$ captures cases where the participant's physical health worsened substantially (by at least one standard deviation) despite expecting it to remain stable or improve. Conversely, if an individual anticipated a deterioration in their health and it indeed worsened, $S_{it} = 0$, since the change was anticipated. This formulation enables us to isolate unexpected negative health shocks, allowing for a more accurate assessment of how unforeseen declines in physical health influence early retirement decisions.

We also include a discrete-time hazard analysis method that helps us observe the influence of health on the timing of the decision to retire early. To achieve this goal, individuals who are at risk of an early retirement at the start of the HILDA survey are examined. This sample is referred to as a stock sample (Jenkins, 1995). We define the stock sample to be comprised of individuals who are aged 50 years or above, have completed a full interview, and are observed to be employed or self-employed in the first wave of the survey. Our main discrete-time hazard analysis sample consists of 9,455 individuals aged between 50 and 64. When additional sociodemographic factors are included, the sample comprises 9,054 individuals, 4,928 males and 4,126 females.

To identify individuals at risk of early retirement in our study, we set an upper age limit of 64, representing the age until which individuals are considered susceptible to early retirement. This choice is consistent with the state retirement age in Australia, which is 65 during the covered period. Individuals aged 64 and below are included in our stock sample, recognising that those over 64 may have different retirement motivations, such as eligibility for social security benefits. Notably, for female participants, we account for variations in the state retirement age until July 2013 when the retirement age was similar for both males and females. This approach determines the observation period for each female participant based on her birth year and the specific state retirement age during the relevant wave.

We follow individuals in every wave until they reach state retirement age or exit the labour market. Individuals cannot be included in all the 18 waves, even if they joined the sample in the first wave of the analysis, since we only estimate the model with workers aged between 50 and 64. Thus, the maximum number of waves per individual is 15. Individuals are followed from the first wave of the survey until they retire, which is assumed to be an absorbing state, are lost to follow-up, or stay in the labour market after state retirement age (65 years old).

We use the death of a close friend as an IV for mental health in both our linear (2SLS) and nonlinear (2SRI) approaches. The relevant question in the HILDA survey asks: "Did any of these happen to you in the past 12 months? Death of a close friend?". We include the IV lagged by one period to align with the lagged structure of our mental health variable. The rationale is that the impact of a close friend's death on an individual's mental health may take time to manifest. Moreover, because the question refers to events in the past 12 months, the death could have occurred at any point within that timeframe, from one month to nearly a year ago, and may already have influenced the previously measured mental health score, depanding on the timing of previous wave. By lagging the IV, we aim to allow sufficient time for the event to affect mental health and better capture its causal impact.

1.5 Results

1.5.1 Summary Statistics

Table 1.4 presents summary statistics for the key variables used in the analysis. Approximately 17% of the observations report being retired, indicating a relatively small but meaningful of early retirement in the sample. The average lagged subjective mental health score is around 74.9 (out of 100), with a comparable mean of 74.4 for the initial mental health measure included in the longitudinal analysis, suggesting general stability in reported mental health over time. However, both variables exhibit considerable variation (standard deviation of roughly 17.4), highlighting notable heterogeneity in mental health within the sample.

The average annual income is AU$86,190, with a substantial standard deviation of AU$81,480, indicating a wide dispersion in earnings. Negative health shocks are relatively rare, affecting only 7.5% of the sample, while 13% experienced the death of a close friend, a variable used to capture exogenous emotional distress. Women constitute approximately half the sample (49%), and 62% of individuals are married. About 23% hold a university degree, and the average household size is 2.5 persons. The local unemployment rate, used as a proxy for labour market conditions, averages around 5.3% but ranges from 1.9% to 8.8%, suggesting some regional variation in economic environments.

Table 1.4 Descriptive Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Retired	44,091	0.169	0.375	0	1
Lagged Mental Health	32,316	74.92	17.42	4	100
Initial Mental Health	38,676	74.41	17.40	0	100
Income	44,091	86.19	81.48	0	1030.59
Negative Health Shock	44,091	0.075	0.26	0	1
Female	44,091	0.49	0.50	0	1
Married	44,091	0.62	0.49	0	1
Degree	41,428	0.23	0.42	0	1
Household size	44,091	2.54	1.23	1	14
Local Unemployment Rate	44,090	5.31	1.11	1.9	8.8
Death of a Close Friend	35,784	0.13	0.34	0	1

Table 1.5 presents descriptive statistics for the variables of interest, disaggregated by individuals' early retirement status and compared with the characteristics of the restricted sample, which includes individuals outside the eligible age range excluded from the main analysis (e.g. pension-aged individuals or younger than 50 years old. On average, individuals report good or very good lagged self-assessed mental health, with a mean of around 75, though this declines among those who have retired early. A notable difference is also observed in the gender composition: approximately 48% of individuals who remain in the labour force are female, compared with around 54% among those who retire early.

Table 1.5 Descriptive Statistics, by Retirement Status and Restricted Sample

Variable	**All**	**Pre-Retirement**	**Retirement**	**Restricted Sample**
Retirement Status	0.169	0	1	0.112
Lagged Mental Health	74.924	75.600	71.913	73.859
Academic Degree	0.233	0.248	0.161	0.228
Income	86.193	92.512	55.214	83.441
Female	0.492	0.483	0.538	0.490
Household Size	2.539	2.623	2.127	3.465
Married	0.619	0.619	0.615	0.297
Local Unemployment Rate	5.314	5.309	5.337	5.292
Negative Health Shock	0.075	0.077	0.067	0.087
Death of a Friend (IV)	0.131	0.128	0.143	0.099

Other variables in the analysis show that most individuals in the sample do not have an academic degree. *Table 1.5* also provides a comparison of the information on the local unemployment rate and the number of individuals in the household. The average local unemployment rate is highest for the post-retirement group at around 5.33%, which is slightly similiar to the pre-retirement group (5.31%). The average household size is 2.54 individuals, with 2.62 individuals for the pre-retirement and slightly lower numbers for post-retirement (2.13). This may indicate that individuals tend to have fewer household members after retirement, possibly due to adult children leaving the household or the passing of a spouse.

The restricted sample provides a useful benchmark to assess whether individuals excluded from the core analysis differ systematically from the main study population. Compared with the full sample, the restricted group displays lower average mental health scores and a higher proportion of negative health shocks, suggesting that younger (<50 years) and older (retirement age) excluded individuals may experience different health dynamics. The restricted group also shows lower marriage rates and substantially larger household sizes, reflecting life-cycle differences in living arrangements. These patterns confirm that the main analytical sample is more homogeneous in terms of age-related socioeconomic characteristics, reinforcing the internal validity of the empirical strategy. There is no evidence of systematic differences between included and excluded individuals that would bias the main results. The restricted sample remains representative of the relevant population, those aged over 50 who are not yet retired and therefore at risk of early retirement.

Interestingly, the number of experiencing negative health shock decreasing post-retirement, suggesting a potential association between an exogenous physical health shock and early retirement. These findings raise intriguing questions about the role of health-related events in influencing retirement decisions and warrant further investigation.

1.5.2 Linear analysis

Table 1.6 presents the LPM estimates of the effect of lagged mental health on the probability of early retirement. Across all three specifications, lagged mental health is statistically significantly and negatively associated with retirement, suggesting that poorer mental health in the previous wave increases the likelihood of early exiting the labour market. In the most basic model (column 1), which includes only the mental health variable, the coefficient is -0.18% and highly statistically significant at the 1% level. This result is robust to the inclusion of sociodemographic controls (column 2), where the coefficient remains negative and significant, though slightly attenuated to -0.15%. When age dummies are also included to account for nonlinear age-related retirement patterns (column 3), the magnitude of the effect increases slightly to -0.17%, reinforcing the conclusion that deteriorating mental health is a strong predictor of early retirement.

Table 1.6 Effect of Lagged Mental Health on Retirement (LPM Models)

	(1)	**(2)**	**(3)**
Dependent variable:		*Retired*	
Lagged mental health	-0.0018^{***}	-0.0015^{***}	-0.0017^{***}
	(0.0001)	(0.0001)	(0.0001)
Sociodemographic controls[a]	–	Y	Y
Age dummies (53–64)	–	–	Y
Observations	32,316	32,057	32,057
R^2	0.0067	0.0495	0.1067

[a] Includes degree, gender, household size, marital status, and regional unemployment rate.

Note: Standard errors in parentheses. $^{***}p < 0.01$

Since mental health is likely endogenous, potentially influenced by unobserved factors that also affect retirement decisions, *Table 1.7* presents estimates from instrumental variable (2SLS) models using the death of a close friend in the previous period as an instrument for mental health. In Column 1, where no control variables are included, the effect of lagged mental health on retirement is substantially larger in magnitude, more than six times greater, compared to the corresponding LPM estimate (-0.18%), and remains statistically significant at the 1% level. This suggests that the LPM model may fail to adequately capture the causal effect of mental health on retirement, potentially due to omitted variable bias or reverse causality.

Table 1.7 Effect of Lagged Mental Health on Retirement (2SLS Estimates)

	(1)	**(2)**	**(3)**
Dependent variable:		*Retired*	
Lagged mental health	-0.0118^{***}	-0.0067^{*}	0.0036
	(0.0032)	(0.0040)	(0.0032)
Sociodemographic controls[a]	–	Y	Y
Age dummies (53–64)	–	–	Y
Observations	29,603	29,432	29,432
First-stage F-stat	59.5	33.55	38.8

[a] Includes degree, gender, household size, marital status, and regional unemployment rate.

Note: Standard errors in parentheses. $^{**}p < 0.05$, $^{***}p < 0.01$. All models use L.ledfr as instrument.

When sociodemographic controls are added (column 2), the coefficient decreases to -0.7% and is statistically significant at the 10% level, indicating the robustness of the relationship after adjusting for confounding variables. However, once age dummies are included in column 3, the coefficient becomes minimal and statistically insignificant.

Across all 2SLS models, the first-stage Cragg-Donald F-statistics (ranging from 33.6 to 59.5) exceed the Stock-Yogo critical threshold of 16.38 for the 10% maximal IV size, indicating that the instrument is sufficiently strong and the estimates are not likely to suffer from weak instrument bias. This exceeds the commonly used rule-of-thumb threshold for identifying a strong instrument, with an F-statistic above 10 and hence indicating that 2SLS may capture the causal effect of mental health on retirement.

Together, these findings support the hypothesis that poor mental health also causally increases the probability of early retirement. However, the attenuation and eventual insignificance of the coefficient once richer controls are included, particularly age, highlight the importance of accounting for life-course factors when modelling retirement behaviour. This highlights the need to apply a longitudinal framework to better capture how individuals' retirement decisions evolve over time. In particular, we expect substantial heterogeneity in retirement behaviour across the age distribution, with older individuals likely exhibiting different patterns and sensitivities compared to their younger counterparts.

The fixed-effects estimation results presented in *Table 1.8* show that the coefficient on lagged mental health is negative, consistent in direction with previous models, suggesting that better mental health is associated with a lower probability of retirement. However, unlike earlier findings, this effect is not statistically significant within the fixed-effects framework. This implies that, after accounting for time-invariant individual characteristics, changes in mental health do not have a statistically meaningful impact on the likelihood of retirement within individuals over time. The lack of significance may be due to limited within-individual variation in mental health or retirement status.

Table 1.8 Effect of Lagged Mental Health on Retirement (Fixed-Effects Model)

Dependent variable:	*Retired*
Lagged mental health	-0.0001
	(0.0002)
Income	-0.0002^{***}
	(0.00004)
Degree	-0.0757
	(0.0408)
Other sociodemographic controls[a]	Y
Age dummies (53–64)	Y
Observations	32,057
Within R^2	0.0688

Note: Robust standard errors clustered at the individual level in parentheses.

$^{***}p < 0.01$

[a] Includes degree, gender, household size, marital status, and regional unemployment rate.

As shown in *Table 1.9*, retirement status remains constant for the majority of the sample across panel waves; only a small proportion of individuals report a change in retirement status, and those who do typically maintain that status over time. This limited within-person variation in retirement partly explains why fixed-effects estimates tend to be less precise or insignificant for retirement outcomes, there is simply insufficient "within" variation to identify the effect reliably.

At the same time, other variables such as household income and household size show statistically significant effects, with higher income and larger household size reducing the

Table 1.9 Within-Individual Variation in Retirement Status

Retired Variation	**Frequency**
0 (No variation)	27,310
1 (Variation)	16,781
Total	44,091

probability of retirement. This indicates that economic resources and family context remain important predictors of retirement decisions.

1.5.3 Discrete-time hazard model

We apply discrete-time hazard analysis to examine how the results of the baseline model change when tracking individuals longitudinally, wave by wave, until they report early retirement or are excluded from the sample due to attrition or other circumstances. This approach allows us to capture the timing of retirement decisions and better account for dynamic factors influencing exit from the labour market.

Figure 1.5 presents Kaplan-Meier survival estimates of the probability of remaining in the labour market across successive time intervals leading up to retirement. Mental health is measured at time $t-1=t_0$ (baseline interview) and treated as a fixed covariate in the survival analysis. The probability of early retirement increases as individuals approach state retirement age. Mental health status is measured at time $t-1$, representing the most recent observation prior to the retirement outcome, to ensure consistency with the econometric specifications presented in the subsequent tables. Individuals reporting excellent or good mental health at $t-1$ are more likely to remain in the labour market at time t compared to those with poorer mental health. Conversely, workers who report very poor mental health exhibit the lowest probability of remaining in the labour market relative to individuals who describe their mental health as excellent, good, fair, or poor.

We estimate a series of discrete-time hazard models using a complementary log-log specification to assess how lagged mental health influences the probability of early retirement across waves. *Table 1.10* presents six model specifications with increasing levels of control. In Model (1), we estimate the baseline relationship between lagged mental health and retirement. The results reveal a strong and statistically significant negative association, with a coefficient of -0.0139 ($p<0.01$). This implies that poorer mental health in the previous wave is associated with a higher hazard of early retirement.

Fig. 1.5 Kaplan–Meier survival estimates, by baseline MH status.

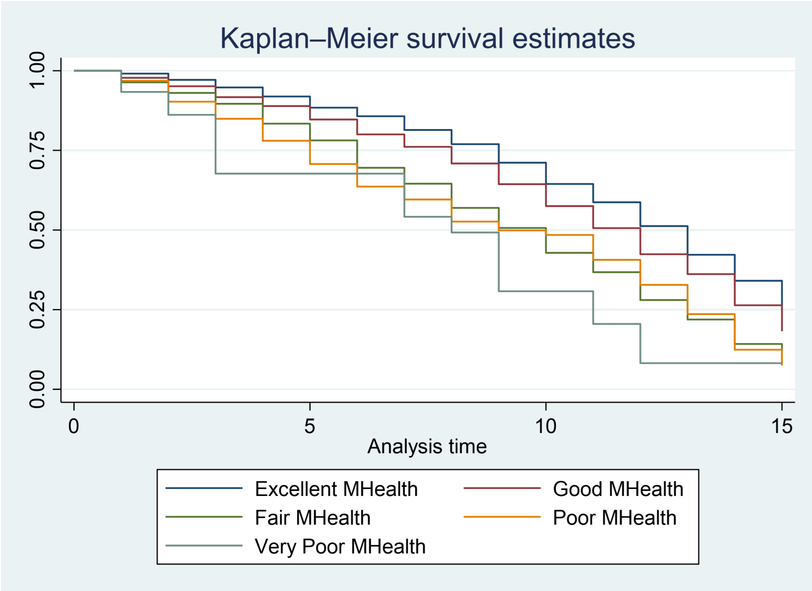

Model (2) introduces a control for initial mental health status, allowing us to net out individuals' initial condition. Both the lagged and baseline mental health variables remain statistically significant, though the magnitude of the lagged effect decreases to -0.09%. This suggests that changes in mental health, rather than static poor levels alone, contribute to early exit from the labour market. Additionally, in Model (3), we add a dummy for experiencing a recent negative health shock. Although the mental health coefficients remain stable and significant, the health shock itself is not statistically significant ($p = 0.178$), implying that acute shocks do not independently predict retirement once mental health is controlled.

Model (4) incorporates household income. The coefficient of lagged mental health remains significant and relatively unchanged -0.086%, while income is also negatively and significantly associated with retirement ($p = 0.017$). This suggests that financial security may partially buffer the impact on retirement decisions. In Model (5), we further control for sociodemographic characteristics including education, gender, household size, and regional unemployment rate. The effect of lagged mental health remains robust and significant (-0.076%, $p < 0.01$), and household size and having an academic degree are also statistically significant.

Finally, Model (6) adds a full set of age dummies to capture non-linear age-specific retirement risks. The coefficient for lagged mental health remains significant and stable (-0.083%, $p < 0.01$), confirming the robustness of the relationship. Notably, the inclusion of age controls substantially improves model fit and reduces the size and significance of

Table 1.10 Effect of Lagged Mental Health on Early Retirement (discrete-time hazard Models)

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	*Early retirement (hazard)*					
Lagged mental health	−0.0139***	−0.0091***	−0.0094***	−0.0086***	−0.0076***	−0.0083***
	(0.0021)	(0.0029)	(0.0029)	(0.0030)	(0.0029)	(0.0029)
Initial mental health	–	Y	Y	Y	Y	Y
Health shock	–	–	Y	Y	Y	Y
Income	–	–	–	Y	Y	Y
Sociodemographics[a]	–	–	–	–	Y	Y
Age dummies (53–64)	–	–	–	–	–	Y
Observations	9,455	9,139	9,139	9,139	9,054	9,054

[a] Includes: degree, gender, household size, marital status, and regional unemployment rate.

Note: Standard errors in parentheses. $^{***}p < 0.01$. Estimated using complementary log-log models.

several other covariates. Many age dummies, particularly those in the early 50s, are large and negative with reduction of the effect size for the older age groups, as expected. This pattern aligns with expectations: the negative association between age and early retirement diminishes as individuals approach the statutory retirement age, reflecting prevailing social norms around the timing of retirement.

We also observe a gradient across educational attainment compared to the baseline category of not having any educational qualifications, with higher level of education are positively linked with a decreasing hazard of retiring. Possessing an academic degree, regardless of other variables, exhibits a substantial positive relationship in the probability of retirement. Moreover, the number of individuals in the household has a negative association with the likelihood to not retire early. According to our results, households with a larger number of living individuals would be less likely to leave the labour market earlier than state retirement age, independently of their health difficulties. The results of the full model are reported in *Table A.3*.

Overall, the discrete-time hazard analysis reinforces earlier findings from the LPM and IV models, demonstrating stronger statistical significance and confirming that lagged mental health is a consistent and significant predictor of early retirement. This effect remains robust even after controlling for initial mental health, physical health shocks, income, sociodemographic characteristics, and age-dummies.

1.5.4 Discrete-time hazard model (including unobserved heterogeneity)

In this section, we include a model with unobserved heterogeneity and test the validation of this model compared to a model without frailty. *Table 1.11* presents the results of a gamma-distributed frailty following the method of Jenkins (1997) for observing frailty. The results show that the expected effect of mental health on early retirement decision is -0.014 for the model without frailty and follows the same pattern of the model with frailty. The effect is statistically significant for both models.

Moreover, we estimate the model fit using information criteria (Akaike and Bayesian information criterions) which were greater with the complementary log-log estimator without frailty compared to the other model. This also shows the similar effect of the independent variable and indicates that our analysis does not set to include frailty. As frailty is not much a concern in this paper and the results do not vary between estimators with or without frailty, we have chosen to apply the complementary log-log regression for the discrete-time hazard analysis assuming no prevalence of frailty in our paper.

Table 1.11 Complementary Log-Log Models of Early Retirement: With and Without Gamma-Distributed Heterogeneity

	No Unobserved Heterogeneity	**Unobserved Heterogeneity**
Lagged mental health	-0.0139^{***}	-0.0138^{***}
	(0.0021)	(0.0036)
Constant	-1.596^{***}	-1.597^{***}
	(0.162)	(0.273)
$\ln(\text{var}(\nu))$	–	-14.852
		(64.841)
Gamma variance	–	3.55×10^{-7}
Observations	9,455	9,455

Note: Standard errors in parentheses. $^{***}p < 0.01$. Unobserved Heterogeneity is estimated assuming Gamma frailty.

1.5.5 Two-stage residual inclusion

To assess the validity of our IV approach, we first estimated the first-stage regression applying LPM, where we regressed mental health score (t-1) on the death of a close friend at time t-2 (*Table A.4*).

The results from the two-stage residual inclusion (2SRI) model presented in *Table 1.12* reinforce the negative and statistically significant impact of lagged mental health on the likelihood of early retirement. Specifically, the coefficient of -1.1% (significant at the 1% level) indicates that a deterioration in mental health in the previous period is associated with a higher hazard of transitioning into early retirement.

Initial mental health status, representing time-invariant baseline psychological conditions, also emerges as a significant predictor, with a coefficient of -0.08% (significant at the 1% level). This highlights that both dynamic (lagged) and baseline (initial) mental health exert meaningful influence on retirement decisions. However, income does not appear to have a statistically significant effect on early retirement in this model, with a small positive coefficient (0.1%) and a p-value above conventional thresholds. This suggests that after accounting for health, sociodemographics, and age effects, income levels may not directly be associated with early retirement decisions in this sample.

Educational attainment, proxied by holding an academic degree, shows a marginally significant negative effect (-1.99%, $p = 0.057$), implying that more educated individuals are less likely to retire early, possibly reflecting greater access to less physically or mentally demanding jobs, stronger labour market attachment, or different preferences regarding continued work.

Table 1.12 Two-Stage Residual Inclusion

Two-Stage Residual Inclusion (2SRI)		
Variable	**Coefficient**	**(SE)**
Second-stage of 2SRI: Early Retirement		
Lagged mental health	-0.011***	(0.003)
Initial mental health	-0.008***	(0.003)
Income	0.001	(0.001)
Degree	-0.199*	(0.105)
Negative health shock	-0.175	(0.194)
Observations	7,335	
First-stage of 2SRI: OLS Regression on Mental Health Status		
Lagged death of a friend	-1.954***	(0.518)
Note: Standard errors in parentheses.		
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$		
F-statistic (First-stage): 14.18		

Other sociodemographic controls, such as household size and local unemployment rate, also exhibit significant associations. A larger household size is associated with a smaller risk of early retirement, while higher local unemployment is positively related to retirement risk, potentially capturing economic discouragement effects. Age dummies behave as expected, with older individuals facing higher retirement hazards. The F-statistic for the instrument in the first-stage regression is 14.2, which exceeds the conventional threshold of 10, suggesting that the instrument is not weak.

These findings highlight the critical role of mental health, both dynamic and baseline, in shaping early retirement behaviour and underline the importance of integrating psychological well-being into labour market and retirement policy frameworks.

1.5.6 Sub-sample: Gender-specific

In this section, we present the results of the sub-sample analysis of estimating males and females separately, as previous research traditionally shows that the results differ when estimating the responses for labour market within the different genders. *Figure 1.6* demonstrate the Kaplan-Meier survival estimates by sex. Based on the gender difference alone, independent of other factors, *Figure 1.6* reveals that male workers in our sample exhibit a higher likelihood of remaining in the labour market compared to females. Additionally, mental health may affect males differently compared to females, with different proportions of individuals with a severe or mild mental health disorder in this sub-groups (Australian Bureau of Statistics, 2022b). The traditional approach when estimating mental health is to observe the effect of mental health on males, however, in this paper we include the results of the two sub-samples (*Table 1.13*).

Fig. 1.6 Kaplan-Meier Survival Estimates, by sex

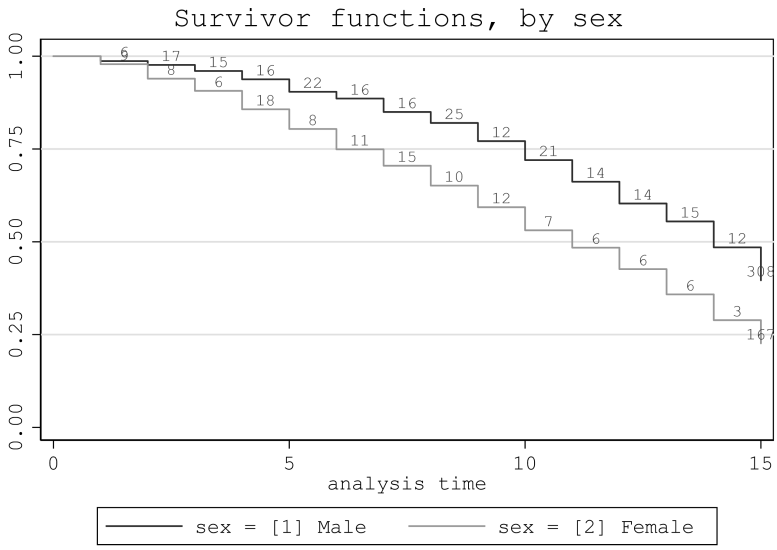

The results in *Table 1.13* present estimates from four different empirical approaches: linear probability models using LPM and 2SLS, and nonlinear discrete-time hazard models using complementary log-log and 2SRI methods across genders. These models aim to evaluate the impact of lagged mental health on early retirement decisions.

Across most models, for both males and females, better lagged mental health is associated with a lower likelihood of early retirement, confirming that mental health issues increase the probability of exiting the labour force prematurely. This effect is consistently negative and statistically significant in nearly all models, except for the 2SLS specification for females, where the sign is positive and statistically insignificant, possibly due to weak instrument bias or sample-specific variation (first-stage F-stats below 10, at 7.907). Notably, the estimated effects tend to be larger in magnitude in the nonlinear models (complementary log-log and 2SRI) for males, particularly after addressing endogeneity in the discrete-time hazard framework.

Comparing across genders reveals notable heterogeneity. The effects of poor mental health on early retirement appear stronger for males than for females in most specifications, underlining the importance of gender-specific analyses when studying retirement behaviour.

Table 1.13 Comparison of Coefficient Estimates Across Models and Gender

	LPM	2SLS	Cloglog	2SRI
Panel A: Female				
Mental health	-0.0014^{***}	0.0116	-0.0057	-0.0063
	(0.0002)	(0.0090)	(0.0041)	(0.0038)
Degree	-0.0307^{***}	-0.0635^{***}	-0.3473^{**}	-0.3685^{**}
	(0.0071)	(0.0210)	(0.1651)	(0.1651)
Income	-0.0006^{***}	-0.0009^{***}	-0.0006	-0.0006
	(0.0001)	(0.0002)	(0.0008)	(0.0008)
Negative health shock	-0.0316^{**}	0.0384	-0.5775^{*}	-0.4792
	(0.0140)	(0.0488)	(0.3400)	(0.3403)
Observations	16,192	14,926	4,126	3,320
F-stat first-stage (2SLS)	–	7.907	–	–
Panel B: Male				
Mental health	-0.0022^{***}	-0.0004	-0.0147^{***}	-0.0164^{***}
	(0.0002)	(0.0034)	(0.0041)	(0.0038)
Degree	-0.0123^{*}	-0.0185^{**}	0.0174	-0.0101
	(0.0065)	(0.0073)	(0.1380)	(0.1383)
Income	-0.0005^{***}	-0.0006^{***}	-0.0010	-0.0012
	(0.0001)	(0.0001)	(0.0008)	(0.0009)
Negative health shock	-0.0078	0.0069	-0.0095	0.0398
	(0.0125)	(0.0214)	(0.2378)	(0.2381)
Observations	15,865	14,506	4,928	4,015
F-stat first-stage (2SLS)	–	39.923	–	–

Notes: Each cell reports the coefficient (standard error in parentheses) for the corresponding variable and model. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Other covariates also exhibit gendered patterns. Education is positively associated with delaying retirement for females, with statistically significant effects across all models. For males, while the direction of the relationship is less consistent and mostly insignificant, the point estimates suggest that having an academic degree may in some cases be linked with earlier retirement, possibly reflecting differences in financial security, job satisfaction, or pension eligibility between educational groups.

Income exhibits consistent effects across linear and nonlinear models and across genders. The association is statistically significant in the linear models for both men and women but not in the nonlinear specifications. This discrepancy may reflect differences in the way time-to-event outcomes are captured or reduced statistical power in the discrete-time hazard models.

Finally, negative health shocks show some evidence of decreasing the risk of early retirement for females, with statistically significant negative coefficients in selected models. However, for males, the effect of health shocks is uniformly insignificant across all specifications. This suggests that physical health shocks may have a stronger influence on early retirement among women, possibly due to differing occupational exposures or caregiving responsibilities.

1.6 Discussion and conclusion

The present study examines the impact of mental health on labour market participation, specifically the decision to remain in the workforce or retire early. Exiting the labour market before reaching the state pension age can have negative implications, including reduced social interaction and financial insecurity. Policymakers therefore need a clearer understanding of how mental health influences early retirement decisions.

Previous literature has explored the relationship between health and labour market outcomes (Disney et al., 2006; Ettner et al., 1997; Frijters et al., 2010; Hamilton et al., 1997; Porru et al., 2019), yet the specific role of mental health among older workers remains underexplored. Empirical challenges such as reverse causality, endogeneity, and measurement error further complicate causal inference. To address these, our study applies a range of econometric methods designed to isolate the causal effect of mental health on early retirement.

First, we validate our derived mental health score using alternative subjective measures, including the Kessler scale (Kessler et al., 1999), and medical diagnoses. The high correlation between these measures suggests that our constructed mental health variable provides a reliable representation of individuals' mental health status.

Second, we estimate linear probability models, including two-stage least squares (2SLS) regressions using the death of a close friend as an instrument for mental health. This event is assumed to affect labour market participation only through its impact on mental health (Frijters et al., 2010).

Although the instrument performs reasonably well, some limitations remain. Only about 13% of respondents reported the death of a close friend, which may restrict variation. Furthermore, the meaning of a "close friend" is subjective and may differ across individuals. Nonetheless, the first-stage F-statistic exceeds 10 for the overall sample, confirming instrument strength. For females, the statistic falls slightly below this threshold at 7.9, but increases to around 40 under the 2SRI specification, reinforcing the validity of our identification strategy.

A further limitation relates to the absence of detailed occupational or industry controls. Job type can plausibly affect physical demands, psychosocial stressors, and exposure to health risks, and may therefore influence both mental health trajectories and early retirement decisions. In our setting, inclusion of occupation may introduce post-treatment bias given that job choice is itself shaped by earlier health and socioeconomic factors. While our models adjust for a broad set of socioeconomic characteristics, including education, income, and household circumstances, that capture some of the systematic differences associated with occupational sorting, we recognise that residual confounding may remain. Future research with richer longitudinal occupational histories would help to further disentangle these channels.

Third, we test for unobserved heterogeneity by estimating discrete-time hazard models both with and without frailty, assuming a gamma mixture distribution (Jenkins, 1995, 1997). The results remain consistent across specifications, suggesting that unobserved frailty is not a major concern.

To account for the timing of early retirement, we use discrete-time hazard models that handle right-censoring and exploit the longitudinal structure of the data. These models show a robust and statistically significant relationship between poorer mental health and higher likelihood of early retirement, even after controlling for physical health and socioeconomic characteristics.

Across all specifications, the results consistently show that individuals with better mental health are significantly less likely to retire early. The estimated effects, although modest in size, are economically meaningful. A one-standard-deviation improvement in mental health reduces the probability of early retirement by roughly 2–4 percentage points. However, an immediate deterioration in mental health would require several unit-level changes in the

SF-36 score to produce a large effect on the probability of exiting the labour market prior retirement age.

Among the alternative empirical specifications considered, the 2SRI model within a survival analysis framework is adopted as the preferred approach. This specification is particularly well suited to the research question, as it combines the advantages of survival analysis with an instrumental variable strategy to address endogeneity. By explicitly modelling time-to-event outcomes, the survival framework allows individuals to be followed over time and appropriately accounts for censoring, thereby exploiting the rich longitudinal structure of the data. In addition, the long-panel nature of the dataset enables a more effective treatment of unobserved heterogeneity at the individual level. The instrumental variable included in this framework demonstrates strong relevance across specifications, reinforcing the credibility of the causal estimates. While alternative models offer complementary insights under different assumptions, the 2SRI survival approach provides a coherent and robust framework that integrates longitudinal dynamics with causal identification, and therefore serves as the primary basis for interpretation of the results.

The gender analysis further reveals that the effect of mental health is stronger and statistically significant for men, while smaller and insignificant for women. This suggests that men's retirement decisions are more sensitive to changes in mental health, possibly reflecting gendered social norms that link male identity more strongly with employment. Women, by contrast, appear more responsive to financial and family-related factors. Lastly, older individuals are more likely to retire early, reflecting both social norms and reduced financial or social penalties for doing so as they approach the state pension age.

The results of our study are broadly consistent with evidence on the effect of general health on labour market outcomes (Bryan et al., 2022; Disney et al., 2006; García-Gómez et al., 2010). However, given the complexity of the relationship between mental health and retirement behaviour, not all studies reach the same conclusion. For example, Andersen et al. (2024) find no causal effect of mental health on early retirement when exploiting a fireworks disaster as an instrumental variable. These contrasting findings highlight the challenges inherent in isolating causal effects in this domain. Our paper contributes meaningfully to the existing literature by providing new evidence on this relationship and by applying a combination of econometric techniques to strengthen identification and robustness.

From a policy perspective, these findings highlight the importance of integrating mental health into retirement and employment policy. Early screening and workplace mental health support can help reduce premature labour market exits. Gender-sensitive interventions, such as promoting open discussions about men's mental health and reducing stigma, could further enhance labour force retention among older workers.

While the estimated effects are relatively modest, they remain economically significant, particularly in the context of cumulative public costs associated with early retirement. Improving access to mental health services and promoting workplace well-being could help delay retirement, thereby extending productive working lives and reducing fiscal pressures linked to ageing populations.

In conclusion, this study contributes new evidence on the causal effect of mental health on early retirement decisions using robust longitudinal and instrumental-variable methods. Our findings underscore that mental health, though not the only driver of early retirement, is a critical factor shaping labour market participation among older workers. Strengthening mental health support, particularly in the workplace, can generate both social and economic benefits.

Chapter 2

The Effect of Absolute and Relative Income on Health: Evidence from Israel

AARON KATZ

Abstract

The relationship between income and health has been extensively studied, yet the association of relative income remains inconclusive. This study extends the analysis by examining the impact of relative income within religious groups, focusing on Jews, Muslims, and Christians in Israel. The analysis incorporates a measure of relative deprivation to compare the effects of absolute and relative income on health outcomes. An ordered probit fixed-effects model with the Mundlak correction is applied as the main specification, while robustness checks using both parametric and semiparametric panel data models are conducted to account for potential misspecification and unobserved heterogeneity. To ensure consistency, the main model is also compared with alternative specifications in which the health outcome is restructured as a binary indicator of good health. The results consistently reveal a significant positive relationship between absolute income and self-reported health across both the main and robustness models, with the exception of one specification. In contrast, relative income, captured through a deprivation measure, exhibits a statistically significant and negative effect on health in all models, suggesting a stronger influence than that of absolute income. Subgroup analyses reveal heterogeneity between religious groups, with relative income showing a positive association with better health for both Jewish and Muslim individuals. Notably, the effect is more pronounced among Muslim participants, indicating that relative income perceptions may vary in salience across cultural or religious contexts.This study highlights the importance of considering relative income and religious group differences in the income-health relationship, particularly in countries with strong religious affiliations. The negative effect of relative income on health suggests that policies addressing income disparities within religious groups could improve overall health outcomes.

Keywords: Relative income, absolute income, health status, religious groups, relative deprivation, income inequality, ordered probit, health disparities.

JEL Codes: I14, I31, D31, C23, Z12

2.1 Introduction

The intricate relationship between income and health has been widely explored across multiple fields, including public health and economics. The World Report on Disability (World Health Organisation, 2021) estimates that 15% of the global population has a health-related disability, with a higher prevalence in low- and middle-income countries, where poverty exacerbates health challenges. Extensive research, drawing on both objective measures (mortality rates, hospital admissions) and subjective indicators (self-reported health), consistently finds that individuals with higher incomes experience better health outcomes (Gibson et al., 2020; Reche et al., 2019; Veenstra and Vanzella-Yang, 2020; Vásquez et al., 2013).

The mechanisms through which income influences health are multifaceted. Income can affect access to education, healthcare, and social capital, creating disparities in health outcomes (Kawachi et al., 1994; Lynch and Kaplan, 1997). In addition, lower-income groups face psychosocial stress from negative social comparisons, further worsening their health (Cui and Chang, 2021; Kawachi et al., 1994). The position of individuals within the income distribution affects not only their material well-being but also their social networks and overall health (Ettner, 1996; Subramanian and Kawachi, 2004).

This paper examines both absolute and relative income channels, building on a growing body of literature that demonstrates how relative income, or one's position within the income distribution, may influence health outcomes independently of absolute income levels (Cui and Chang, 2021; Jones and Wildman, 2008; Kawachi et al., 1994). Although absolute income has been extensively studied (Ettner, 1996; Subramanian and Kawachi, 2004), the relative income hypothesis has been comparatively underexplored (Fiscella and Franks, 1997; Kennedy et al., 1996). Relative deprivation can result in poorer health even for individuals with similar absolute income levels.

Individuals who view themselves as worse off than their peers may experience psychosocial stress and pressure, feelings of inferiority, and reduced social cohesion, which can negatively affect both mental and physical health through behavioural and physiological pathways (Jones and Wildman, 2008; Kawachi et al., 1994; Wilkinson, 1996). This framework forms the basis of the relative income hypothesis, which argues that health is shaped not only by what individuals have, but also by how their income compares to that of others.

In addition to income, this study explores the role of religion in shaping health outcomes. In Israel, where religion plays a central role in shaping societal norms and behaviours, the intersection of religion and health is particularly significant (Abu Riha, 2015; Cooperman et al., 2016; Holt et al., 2009). The religious diversity in Israel, with Judaism, Islam, and

Christianity as the major faiths, provides a unique context to investigate how religious beliefs, practices, and dietary customs impact health. This study takes advantage of Israel's religious heterogeneity to conduct in-depth analyses of how religion moderates the relationship between income and health.

This research makes a substantial contribution by filling several important gaps in the existing literature. First, it applies both absolute and relative income hypotheses to a unique Middle Eastern context, including panel data from Israel's Central Bureau of Statistics (CBS), which provides an under-explored yet rich dataset. The diverse Israeli population, with its pronounced religious and socioeconomic diversity, allows for a nuanced analysis that has not been previously undertaken. Second, this paper applies both parametric and semi-parametric approaches to account for an unspecified functional form of income on health, allowing us to compare how this non-linearity impacts the outcome relative to the parametric approach. Third, incorporating a binary health outcome variable in an additional analysis, alongside the self-reported health status already included within the main model, enhances the accuracy of health measurement by reducing potential biases in self-reported data. Finally, this paper goes beyond previous research by comparing income-health relationships across different religious groups, providing critical insights into how religious affiliation shapes health outcomes in a multi-faith society.

The findings reveal a consistent positive association between absolute income and self-reported health across both primary specifications and robustness checks, with the exception of one model. Relative income, captured through a deprivation measure, shows a consistently negative and statistically significant effect on health across all models, an effect that is notably stronger than that of absolute income. To explore potential non-linearities in the income-health relationship, a semiparametric approach was included as a validation exercise. This method confirmed the robustness of the relative income effect, which also exhibited a statistically significant marginal impact compared to other models.

Heterogeneity analysis by religious affiliation revealed that the positive effect of absolute income followed a similar pattern for Muslim and Christian individuals. However, when relative deprivation was accounted for, the patterns shifted, with Muslims and Jews displaying more comparable trends. These results highlight the importance of considering both relative income and religious group differences when analysing the income-health relationship, particularly in settings where religious identity plays a prominent social role. The findings suggest that addressing income disparities within religious communities could have meaningful implications for improving population health.

This study not only advances theoretical understanding, but also has significant practical implications for policymakers. By highlighting the differential impact of income and relative deprivation in general and across religious groups, it offers new insights for designing targeted, culturally-sensitive interventions to reduce health disparities. The findings show the importance of addressing both absolute and relative income inequalities in public health policy, rather than focusing solely on absolute income, particularly in religiously diverse societies such as Israel.

The paper is structured as follows. Section 2.2 reviews the existing literature on income and health. Section 2.3 outlines the methodology, including the relative deprivation measure. Section 2.4 describes the data source, and Section 2.5 presents the results of the main- and sub-sample analyses. Section 2.6 discusses the implications of these findings, while Section 2.7 concludes the paper.

2.2 The literature

This paper investigates the relationship of absolute and relative income with health. Specifically, it studies the hypothesis that an individual's rank in the income distribution influences their self-reported health status. This section reviews previous empirical studies investigating these hypotheses to provide a context for this paper.

2.2.1 Absolute income hypothesis

The absolute income hypothesis claims a positive relationship between income and health status (Ettner, 1996; Ruckert et al., 2017; Subramanian and Kawachi, 2004). Numerous studies have extensively explored the link between absolute income and self-reported health, consistently revealing a positive relationship. Although the precise impact may vary depending on the context and population under study, most of these inquiries suggest that higher income correlates with improved health outcomes. Notably, Ettner (1996) included various instrumental variables (e.g. the state unemployment rate, work experience, parental education, and spouse characteristics) to assess the influence of income on physical and mental health and other health-related indicators, finding a statistically significant positive association between income and better health. These findings align with the broader body of research on the absolute income hypothesis, which implies that higher income may be a protective factor for health.

Previous empirical studies across various fields, including epidemiology, social welfare, economics, and sociology, have consistently demonstrated a significant non-linear relationship between income and health when using individual-level data (Adams et al., 2003; Feinstein, 1993; Smith, 1999). Preston (1975) used macro relationships at a country level to observe the effect of income on mortality rates. The results showed a concave relationship between income (as part of economic development) and health, where increases in income had a greater effect on mortality rates in lower-income countries than in higher-income countries. This indicates that there might be a non-linear relationship between income and health at an individual level. To address the non-linear association between income and health in the parametric modelling in this study and following previous research (Ecob and Smith, 1999; Ettner, 1996), the logarithm of income is included. A semiparametric approach is also applied as a robustness check, assuming a flexible functional relationship between income and health.

Furthermore, this paper is included with the Hausman-Taylor (HT) estimator, introduced by Hausman and Taylor (1981), which provides a useful approach for estimating panel data models where some explanatory variables are correlated with individual-specific effects, while others are not. Unlike standard fixed-effects models, which eliminate all time-invariant variables through differencing, the HT estimator allows for the inclusion of time-invariant covariates by exploiting the exogeneity of a subset of the variables as instruments. Specifically, it distinguishes between time-varying and time-invariant variables and between those that are endogenous or exogenous with respect to the unobserved individual effect. This feature makes an additional contribution of the effects of important time-invariant characteristics, such as education, gender, while also addressing potential endogeneity in the model. In the context of health economics, the HT approach has been widely used to study income-health relationships, healthcare utilisation, and employment outcomes, where both observed and unobserved heterogeneity play a central role (Contoyannis et al., 2004; Jones, 2009). Its flexibility provides a valuable compromise between fixed-effects and random-effects estimators, making it particularly well-suited for analysing complex, rich panel datasets such as this used in this study.

This study adopts a similar approach to prior empirical research and considers the impact of absolute income on health including it in unique Israeli data and various econometric methods. Although we do not expect to find significant deviations from the existing literature, we include these findings to compare them with outcomes derived from exploring the relative income hypothesis.

2.2.2 Relative income hypothesis

The relative income hypothesis posits that the impact of income on health is not solely determined by the level of an individual's income, but also by their position in the income distribution relative to others. This hypothesis was first proposed by Duesenberry (1949) to explain individual saving behaviour in the United States, through the development of a consumption function that incorporated the current income of others. This approach aimed to reconcile the differences observed between cross-sectional and time-series characteristics of consumption data. Empirical results showed that the percentage of savings is positively correlated with the percentile position in the income distribution, but is independent of absolute income (Cai et al., 2022). This suggests that relative income can have a greater impact on savings behaviour than absolute income.

Similarly to the absolute income hypothesis, the relative income hypothesis has been empirically tested across various disciplines, including economics, sociology, and public health, in relation to health outcomes (Deaton, 2001; Kawachi et al., 1994; Wilkinson, 1996). Research in social psychology further supports this hypothesis by showing that individuals in different income groups experience varying levels of social pressure and comparison-based stress, which can influence their health and wellbeing (Festinger, 1954; Macleod and Smith, 2003). For example, individuals in low-income groups experience greater psychological pressure and a lack of control, which can have an impact on their mental and physical health. This pressure can be influenced by factors such as the state unemployment rate, work experience, parental education, and spousal characteristics (Ellaway et al., 2012; Frank, 1985).

Studies also suggest that relative deprivation compared to others in a group may cause additional stress and depression, which may affect health status via health-compromising behaviours, such as smoking and increased consumption of alcoholic beverages (Cohen et al., 1997). The effects on health can include an increase in work hours, levels of unhappiness, absence from work, the risk of attempted suicide, and premature death (Clark and Oswald, 1996; Ellaway et al., 2012; Frank, 1985; Jones and Wildman, 2008). Thus, the positive association between health and an affluent lifestyle highlights the importance of understanding the association between relative income and health (Kennedy et al., 1996).

Constructing a deprivation measure to estimate the relative income hypothesis necessitates careful consideration of the reference group. Various reference groups have been used in previous studies. The most commonly used method is a national-level reference group, which assumes that the whole population of a country serves as the comparator group (Bygren, 2004; Deaton, 2001; Jones and Wildman, 2008; Kaplan et al., 1996; Lindley and Lorgelly, 2005; Wilkinson, 1996). The existing literature offers different perspectives on defining reference groups. Runciman (1966) proposed two potential reference groups: individuals in the same social situation, such as those within the same social class or religion, and individuals who reside nearby. Bygren (2004) found that workers in Sweden generally compare themselves with the whole population, which aligns with the view of Deaton (2001). Wilkinson (1996) examined country-level inequality measures, whereas Blanchflower and Oswald (2004) applied relative income measures based on individual income divided by state income.

The results of previous empirical studies have varied depending on the reference group. Some studies have used the mean income of a community as an indicator of relative income but did not find correlation with the outcome variable (Robert, 1998). Another study by Eibner and Evans (2005) found that relative deprivation significantly contributes to health-compromising behaviours. However, Jones and Wildman (2008) conducted a study that supported the negative relationship between relative deprivation and health by applying parametric and semiparametric methods to allow flexibility in specifying the effect of income on health. Jones and Wildman (2008) applied a national reference group, where each individual compares their income rank to every other individual in the sample. Our paper adopts the relative deprivation approach introduced by Jones and Wildman (2008) and Clark (2003), comparing the income of each participant within the dataset to every other individual within the same year and comparing the results across the different religious groups.

2.2.3 Religious Groups

Religion exerts a profound influence on individual behaviours and societal norms, particularly regarding attitudes towards alcohol, drug consumption, and dietary habits, all of which can significantly impact health outcomes. In many religious traditions, such as Islam, the consumption of alcohol is strictly prohibited due to religious teachings emphasising the importance of maintaining physical and spiritual purity and avoiding behaviours that may lead to harm or addiction. In contrast, in Judaism, the consumption of wine is not only allowed but encouraged for adults, often serving as an integral component of religious rituals and festive meals. The practice of reciting a blessing over wine, known as "Kiddush", is a central element in sanctifying religious holidays. These different attitudes towards alcohol consumption reflect the diverse teachings and customs of religious communities (Fuller, 1996).

Numerous studies have found that religious involvement is associated with better health outcomes. For example, a study by Idler et al. (2003) found that religious attendance was associated with better self-rated health and lower levels of depression and anxiety. There are several possible explanations for the link between religion and health. One is that religion provides individuals with a sense of purpose and meaning, which can lead to greater resilience in the face of stress and adversity (Koenig, 2012). Furthermore, religious involvement often involves social support and a sense of community, which can buffer against the negative effects of stress (Levin, 2013). Finally, some research suggests that religious involvement may be associated with healthier lifestyles, such as lower rates of smoking and alcohol consumption (Holt et al., 2009).

While the evidence linking religious involvement and health outcomes is robust, there are some reservations to consider. For example, the link may not be causal, as religious individuals may be more likely to engage in healthy behaviours and seek medical care when necessary (Koenig, 2012). Furthermore, the relationship may vary depending on the specific religious group or denomination and the individual's cultural context. Levin (2013) examines the relationship between religious behaviour, health and well-being among Jewish individuals in Europe using data from the European Social Survey. The study found that religious behaviour was positively associated with better health and well-being outcomes, including lower rates of chronic illness and depression. The findings suggest that practicing a religion can promote better health outcomes, and this relationship should be considered when examining the impact of income on health in different religious groups.

Religion is a culturally significant factor in Israel that may influence health-related behaviours, including daily routines, attitudes toward various issues, political views, and levels of community support. According to a report by the Pew Research Center and Israel's CBS, religion plays a crucial role in the lives of individuals belonging to the three dominant religious groups in Israel (Cooperman et al., 2016). The report highlights significant differences between various religious communities and emphasises a strong sense of attachment to one's religion of birth, with approximately 99% of the participants reporting the same religion as their birth religion. Additionally, about 80% of the respondents reported that they would not be pleased if their children married someone of a different religious background.

This study seeks to build on this research by considering the three religious societies to capture additional variations that may arise from belonging to a specific religious group. Specifically, the study's model will be applied to these sub-samples to investigate the heterogeneous effects of income on health across these groups.

2.3 Methods

This paper adopts a model that examines the association between self-reported health and absolute income, while also incorporating a relative deprivation measure to compare the relative income effect. This allows us to better capture the direct effect of absolute income on health and abstract from any effect that the relative position in the income distribution might have.

We use panel data and incorporate the Mundlak approach to account for fixed-effect variation. Since non-linear models such as those required for self-assessed health do not support fixed-effects due to incidental parameter problem, we apply the Mundlak approach (Mundlak, 1978). This involves calculating the mean for each participant in our dataset based on the time-varying variables. When the deprivation measure is excluded, we include the mean of of the lag of logged income per individual as an additional variable. Taking the log to deal with non-linearities, and secondly to lag to deal with simultaneity. However, when the model is applied with relative income, both the mean of lagged logged income and the deprivation measure for each participant are included.

2.3.1 The absolute income hypothesis

To capture the categorical self-reported health, we use an ordered probit model. It is important to note that our model can be viewed as a reduced form specification, as it does not include objects of choice, such as the quality of medical care or health supply.

We use a Mundlak formulation (Mundlak, 1978) as the main model to control for correlation between the potentially exogenous variables and the time-invariant error term. In random-effects models, if there is a correlation between the control variables and the error term (u_i), the estimated coefficients will be biased and inconsistent. The reduced-form model with Mundlak approach can be summarised as follows:

$$H_{it}^* = \beta_0 + I_{it}\Gamma + X_{it}\theta + u_i + e_{it} \quad (i = 1, \ldots, N; \quad t = 1, \ldots, T_i) \tag{2.1}$$

Where, H_{it}^*, is the latent outcome with the categorical self-reported health status of individual i at time t. I_{it}, is the reported lag of logged income. X_{it}, is the vector of other socioeconomic factors (such as, marital status) that affect our outcome variable, H_{it}^*. u_i is the individual specific and time-invariant random error component, which following Mundlak approach (Mundlak, 1978) is assumed to be a function of the individual-specific means of, I_{it}, and assumed to be drawn from a normal distribution ($u_i = \bar{I}_i\Gamma_m + \varepsilon_i$), e_{it} is the time and individual-specific error which is assumed to have a standard normal distribution uncorrelated with I_{it}, X_{it} and u_i, across individuals and waves.

As we do not have a natural scale for the latent variable, the variance of the idiosyncratic error term is restricted to one. In our data, the latent outcome H_{it}^* is not directly observed. Instead, we observe an indicator corresponding to the category in which the latent variable falls (h_{it}). The observation mechanism can be expressed as:

$$H_{it} = j \text{ if } \mu_{j-1} < H_{it}^* \leq \mu_j, \quad j = 1, \ldots, z \tag{2.2}$$

Here, $\mu = -\infty$, $\mu_j \leq \mu_{j+1}$, and $\mu_z = \infty$. Assuming that the error term has a standard normal distribution, the probability of observing a specific category of the health status reported by individual i at time t (H_{it}), conditional on the regressors and the individual effect, is given by:

$$P_{itj} = P(H_{it} = j) = \Phi(\mu_j - \beta_0 - I_{it}\Gamma - X_{it}\theta - \bar{I}_i\Gamma_m - \varepsilon_i) - \Phi(\mu_{j-1} - \beta_0 - I_{it}\Gamma - X_{it}\theta - \bar{I}_i\Gamma_m - \varepsilon_i) \tag{2.3}$$

Here, $\Phi(\cdot)$ represents the standard normal distribution function. This formulation highlights that it is not possible to separately identify an intercept in the linear index (β_0) and the cut-off points (μ), where $u_i = \bar{I}_i\Gamma_m + \varepsilon_i$; the model only identifies ($\mu_j - \beta_0$). To address this, we have adopted a conventional normalisation by setting $\beta_0 = 0$.

2.3.2 The relative income hypothesis

It is acknowledged that individuals' perceptions of their relative socio-economic situations may occur at different levels of awareness or without conscious deliberation. The approach of Eibner and Evans (2005), subsequently applied by Jones and Wildman (2008) to construct the relative income hypothesis, is included in our paper and is presented through the following specification:

$$H_{it}^* = \beta_0 + \beta_1 RD_{it} + I_{it}\Gamma + X_{it}\theta + \bar{I}_i\Gamma_m + \overline{RD}_i\beta_{1m} + \varepsilon_i + e_{it} \tag{2.4}$$

Where RD_{it} indicates the relative deprivation measure, and the time-invariant random error component, u_i, additionally includes the individual mean of, RD_{it}. Here, $u_i = \bar{I}_i\Gamma_m + \overline{RD}_i\beta_{1m} + \varepsilon_i$.

We estimate RD of individual i, at time t, comparing their income, I_{it}, to another individual with income, m, following the deprivation measure of Hey and Lambert (1980), as follows:

$$D(I;m) = \begin{cases} m - I & \text{if } I < m \\ 0 & \text{if } I \geq m \end{cases} \tag{2.5}$$

Equation(2.5) presents a measure of relative deprivation, facilitating pairwise comparisons between individuals. We anticipate that individuals with lower incomes relative to others may experience a heightened sense of deprivation. To incorporate this measure at the individual level, we divide the aggregate values into weighted groups of individuals with incomes below and above $'m'$. Following the methodology proposed by Chakravarty (1990) and Jones and Wildman (2008), we specify the individual-level relative deprivation measure as follows:

$$RD = d_I(F) = \mu[1 - F_1(I)] - I[1 - F(I)] \tag{2.6}$$

Equation(2.6) was applied to introduce a measure of relative deprivation for each individual in our sample. The equation incorporates key variables, including mean income (μ), ($F_1(I)$) the cumulative proportion of total income at income level I, and ($F(I)$) the cumulative proportion of the population up to the individual with income level, I. By applying Equation (2.6) to all individuals included in our analysis, we can effectively account for relative deprivation in our model.

2.3.3 Robustness check

An alternative outcome measure was constructed as a binary indicator reflecting whether individuals are classified as being in good health. This approach allows us to mitigate against potential measurement error of an ordinal self-assessed health status, in addition to exploring additional models, such as the Hausman-Taylor estimator, and via a linear probability model to provide a more comprehensive analysis. Furthermore, by assuming that income affects good health through an unspecified functional form, we can additionally apply a semiparametric approach.

A Probit model with random-effects is included alongside the Probit Mundlak approach for fixed-effects, used as the primary method in our main model. This allows for a thorough comparison between the two models. The key difference between these models lies in their assumptions about the correlation between individual-specific effects (u_i) and the set of regressors, (X_{it}). In the random-effects model, the u_i are treated as random draws from a normal distribution, implying that the individual effects are independent of the covariates, allowing for inferences about population characteristics. On the other hand, the fixed-effects estimator assumes a conditional relationship upon the individual-specific effects, u_i, permitting correlation between u_i and the regressors, (X_{it}).

There are two significant limitations with these models. First, the user must decide entirely on whether they believe that there is correlation between u_i and the regressors, (X_{it}), or not (i.e., whether to use the fixed-effects estimator or the random-effects estimator). Second, if it is more plausible that the individual effects are related to the regressors, estimating time-invariant explanatory variables becomes impossible with the fixed-effects estimator, which requires 'within' variation (i.e., variation over time within individuals). To overcome these limitations, we incorporate the Hausman-Taylor (HT) model (Hausman and Taylor, 1981). The HT estimator uses both between-group and within-group variations in exogenous variables as instruments, creating a hybrid structure of fixed- and random-effects, as detailed below.

The semiparametric approach addresses the challenge posed by the non-linearity in the relationship between health and income. If this relationship is incorrectly specified, the neglected non-linearity can act as an omitted variable, leading to biased estimates if it is correlated with other unobserved variables. To address this, a semiparametric method is applied. By not specifying the exact nature of the relationship between health and income, this method provides a more flexible approach to obtain an unbiased coefficient for the deprivation variable and allows us to explore the shape of the health-income relationship.

Including these additional models offers clarity on the effect of relative deprivation (relative income hypothesis) on health and a consistency check with the results of the main model.

Hausman-Taylor

The Hausman-Taylor model is an extension of random-effects panel data models, designed to handle situations where some covariates are correlated with the unobserved individual-level random effect. The HT estimator leverages the mean values of time-varying exogenous variables to identify parameters linked with time-invariant endogenous variables, potentially over-identifying them. This estimator is as accurate as the within estimator and can avoid potential inconsistency issues.

When using random effects, any correlation between exogenous variables and the error term (u_i) can lead to biased and inconsistent coefficient estimates. To detect such correlations, a Hausman test (Hausman and Taylor, 1981) is applied. This test compares the estimates from the fixed- and random-effects models. If the null hypothesis of correct specification (exogeneity) holds, the random effects estimator's results should closely match those from the fixed effects method for time-varying regressors. The test statistic is computed as $(\hat{\beta}_{FE} - \hat{\beta}_{RE})/\sqrt{[se(\hat{\beta}_{FE})]^2 - [se(\hat{\beta}_{RE})]^2}$. The HT estimator includes exogenous time-varying and time-invariant variables as instrumental variables, as they are not correlated with the error term (ε_{it}), enabling the estimation of our model, which includes potentially endogenous time-varying variables: lagged-log income and relative deprivation rank.

The Hausman–Taylor approach estimates a model of the following form:

$$GH_{it} = X_{1it}\beta_1 + X_{2it}\beta_2 + X_{3i}\beta_3 + u_i + \varepsilon_{it} \tag{2.7}$$

Here, let GH_{it} represent good health of individual i at time t, while X_{1it} represent a $1 \times k_1$ vector of observations on exogenous, time-varying age dummies. These are assumed to be uncorrelated with u_i and ε_{it}. X_{2it} is a $1 \times k_2$ vector of observations on endogenous, time-varying variables- income and relative deprivation (when included). These are assumed to be (possibly) correlated with u_i, and include the within-transformed variables ($\tilde{X}_{2it} = X_{2it} - \overline{X}_{2i}$), as instruments. X_{3i} is a $1 \times g_1$ vector of observations on exogenous, time-invariant variables- marital status, highest educational attainment, main language spoken, and gender, assumed to be uncorrelated with u_i. Considering the nature of our dataset, marital status, highest educational attainment, and main language show little to no change within each individual and are presumed to be constant over time. Finally, ε_{it} is the time and individual specific, and is assumed to be uncorrelated with X_{it} and u_i.

Semiparametric models

Parametric estimation may lead to biased coefficients if the relationship with income and health is nonlinear. Semiparametric or partially linear estimation, as proposed by Robinson (1988), addresses this issue by allowing for an unspecified relationship between health and income.

We specify a partially linear model as follows:

$$GH_{it} = g(I_{it}) + X_{it}\theta + u_i + \varepsilon_{it} \tag{2.8}$$

In this model, $g(I_{it})$ represents the unknown function of income. Using the method described by Robinson (1988), we transform the model by taking conditional expectations and then subtracting these expectations from Equation(2.8). This yields:

$$H_{it} - \mathbb{E}[H_{it}|I_{it}] = (X_{it} - \mathbb{E}[X_{it}|I_{it}])\theta + \nu_{it} \tag{2.9}$$

The expectation $\mathbb{E}[X_{it}|I_{it}]$ is constructed using a semi-parametric approach. Specifically, we assume that the conditional expectation function can be decomposed into a parametric part and a non-parametric part. We use a linear component for the parametric part, capturing the main effects of the covariates, and a non-parametric component to model the potential non-linearities of income. Estimating Equation (2.9) provides root-N consistent estimators of the β coefficients (Robinson, 1988). It is possible to estimate θ consistently without modelling g(I) explicitly. The unknown function of income, g(I), can be estimated by regressing $(H_{it} - X_{it}\hat{\theta})$ on I_{it} nonparametrically. Where, v_{it}, represents the transformed error term, which accounts for the conditional expectation taken with respect to I_{it}.

2.4 Data

This study uses longitudinal data from Israel's Central Bureau of Statistics (CBS), covering seven waves between 2012 and 2019, including approximately 4,500 households and 14,000 adults. Initially conducted face-to-face, interviews were conducted by telephone from the fourth wave onwards. The survey collected health status and sociodemographic data, including religious affiliation. Only individuals participating in at least two consecutive waves were included to support longitudinal analysis.

The analysis focuses on individuals likely to be active in the labour market, ensuring the inclusion of participants with reported income who are employed or self-employed. This approach excludes individuals ($n = 30$) who are not part of the active workforce, such as those who are severely ill or serving in the military, as they receive a fixed basic income, thereby offering a more stable and relevant basis for analysing the relationship between income and health outcomes.

The initial sample for this study consisted of 114,913 observations. A total of 30 observations were excluded due to missing employment data. Additionally, 55,312 observations were excluded based on age criteria: 31,111 observations were removed for being underage, and 24,201 observations were removed for being overage. This left 17,030 observations that were not available in at least one consecutive wave of the data, further reducing the sample size. After applying these filters, the remaining sample included a total of 42,541 observations for analysis.

These exclusions may have implications for the representativeness of the estimation sample. In particular, attrition and missing data are unlikely to be random and may be correlated with labour market attachment or health status, potentially leading to a sample that is more stable, healthier, or more strongly attached to the labour force than the underlying population. As a result, the estimated effects should be interpreted as applying to individuals with relatively consistent panel participation, and the findings may understate the magnitude of relationships for more transient or disadvantaged groups.

Fig. 2.1 Sample construction and exclusion criteria.

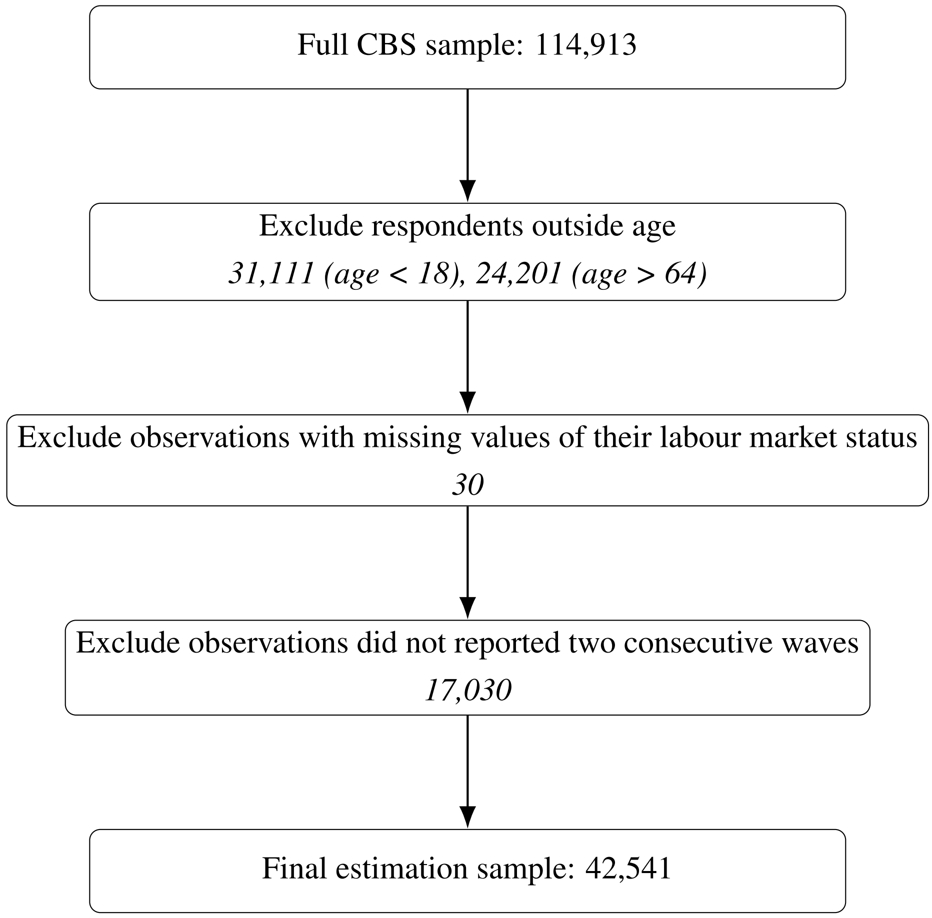

2.4.1 Dependent Variable

The primary outcome is self-assessed health, collected via a general health questionnaire, similar to surveys like HILDA and BHPS. Participants rated their health on a scale of one to four, which was adjusted so that higher values represent better health. To enhance robustness and apply other models, the health status was additionally converted into a binary variable, where "1" represents good or excellent health and "0" represents fair or poor health.

2.4.2 Independent Variables

Income: The income variable, drawn from CBS data. Income was deflated using the consumer price index (base year: 2012) and calculated annually. A lagged income variable was used to account for delayed health impacts and to measure income changes over time. The natural logarithm of income was used to capture its non-linear relationship with health, as in previous studies (Ettner, 1996).

It is important to note that our analysis focuses on individual income rather than equivalised household income. This choice has implications for the interpretation of our findings. Individual income provides a direct measure of a person's own earnings capacity and labour market attachment, thereby capturing the psychosocial dimensions of economic position, such as autonomy, self-esteem, and perceived social status, that may influence health outcomes. In contrast, equivalised household income reflects shared material resources and living standards within the household, accounting for economies of scale and intra-household resource sharing. While the latter offers a broader perspective on material well-being, it may obscure variations in personal economic circumstances, particularly in households with unequal income contributions. Consequently, focusing on individual income allows for a clearer examination of how personal economic standing relates to health, independent of household-level effects.

Relative Deprivation: Relative deprivation was constructed following the approaches of Hey and Lambert (1980), Chakravarty (1990), and Jones and Wildman (2008). The measure captures the extent to which an individual's income falls short relative to others in their reference group. For each wave, we calculated relative deprivation within both the full national sample and within sub-populations defined by religious affiliation to account for potential differences in social comparison groups.

Individual income was expressed in constant New Israeli Shekels(NIS). The relative deprivation index was derived by comparing each individual's income to the income distribution of others in the sample, where deprivation arises only when an individual's income is lower than that of others. The index therefore reflects both the intensity and frequency of income shortfalls relative to peers.

To facilitate interpretation, higher positive values indicate greater relative deprivation, corresponding to a lower position in the income distribution. In contrast, absolute income, measured in thousands of NIS captures the direct level of material resources. Including both absolute and relative income in the analysis allows us to distinguish between the effects of material well-being and perceived social standing on health outcomes.

Other Variables: The analysis also includes controls for age, sex, education, marital status, and primary language (Hebrew, Arabic, or English), following the literature on individual-level data.

2.5 Results

2.5.1 Summary Statistics

Table 2.1 presents the mean and standard deviation of the variables used in this study, along with their minimum and maximum values. The study includes binary variables for *Hebrew, Arabic and English* languages. Most of the sample speaks *Hebrew* as their main language, accounting for approximately 67%, followed by *Arabic* with around 21%. Only 9.8% of the individuals in the sample report *English* as their main language of communication. The variable *female* indicates the proportion of female participants, with a mean of 40%, while the variable *married* measures whether the individuals are married, with a mean of nearly 70%. The Self-Assessed Health (*SAH*) represents the rescaled individuals' reported general health condition, ranging from 1 (Poor) to 4 (Excellent), with an average of 3.25, indicating a good health status for most participants in the data. The variable of lagged-log income (*loglincome*) measures the natural logarithm of lagged income, with a mean of 11.247, ranging from 0.006 to 14.27. The age categories range from 21 to 59, with ten-year intervals (except for the first age group of 21 to 29, as it is expected that individuals will remain in the army until the age of 21). Most of the participants in the sample identified as Jewish, comprising 66.6%, followed by Muslims, who constituted 16% of the sample, and Christians representing only 3.6%. The mean and standard deviation are provided for each age category.

Figure 2.2 and *Figure 2.3* depict the distribution of the mean frequencies of reported self-assessed health status among religious groups in the first and second waves of the analysis, respectively. Notably, in the first wave, Christian participants reported more instances of good health than excellent, unlike the other religions, where excellent health was the most commonly reported. In the second wave, Muslim participants reported good health almost as frequently as fair health, while for Christian participants, excellent health was the most reported, similar to other religious groups. Muslim participants have the greatest ratio of poor health reported compared to other religions. For Jewish participants, the pattern observed in the first wave persisted in the second wave. As hypothesised in the earlier discussion, these variations in the frequencies of reported health status across religious groups suggest that religious affiliation influences health status.

Table 2.1 Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max	Additional Information
wave	42,541	4.358	2.636	1	7	
loglincome	20,846	11.247	1.189	0.006	14.274	
						Between: Std. Dev. = 1.109
						Within: Std. Dev. = 0.653
dep	29,113	0.211	0.907	-8.881	1.597	
						Between: Std. Dev. = 7.927
						Within: Std. Dev. = 4.411
SAH	42,422	3.253	0.890	1	4	
hebrew	42,541	0.668	0.471	0	1	
arabic	42,541	0.208	0.406	0	1	
english	42,541	0.098	0.298	0	1	
female	42,541	0.401	0.490	0	1	
academicdegree	42,541	0.365	0.481	0	1	
married	42,541	0.691	0.462	0	1	
age2129	42,541	0.110	0.313	0	1	
age3039	42,541	0.198	0.398	0	1	
age4049	42,541	0.192	0.394	0	1	
age5059	42,541	0.161	0.368	0	1	
jewish	42,541	0.666	0.472	0	1	
christian	42,541	0.036	0.186	0	1	
muslim	42,541	0.162	0.368	0	1	
mean_loglincome	32,067	11.134	1.087	0.006	13.986	
mean_dep	32,609	0.247	0.806	-6.205	1.597	

2.5.2 Absolute and Relative Income Hypotheses

This section presents the marginal effects of applying the ordered probit model to Equation(2.4). Initially, the study investigates the absolute income hypothesis by exclusively including the income variable of interest, as illustrated in *Table 2.2*. Subsequently, the analysis encompasses the deprivation measure. The discussion mainly revolves around the relationship between absolute and relative income on health. Recognising the significance and growing research interest in the relationship between education and health, the study included this variable in the marginal effect table. The main analysis results and additional sociodemographic variables that may be relevant based on other research in this area can be found in *Table B.1*.

The coefficients reported in *Table 2.2* present the marginal effects estimated using the Mundlak specification of the ordered probit model (Mundlak, 1978). These estimates capture the percentage-point change in the probability of reporting each health category associated

Fig. 2.2 Self-reported Health, by Religious Groups in wave 1

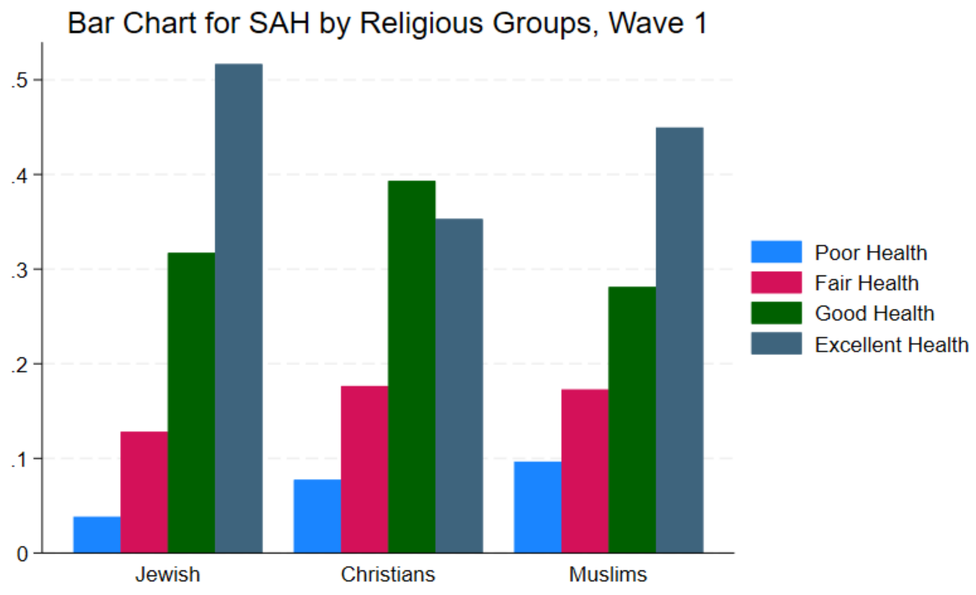

Fig. 2.3 Self-reported Health, by Religious Groups in wave 2

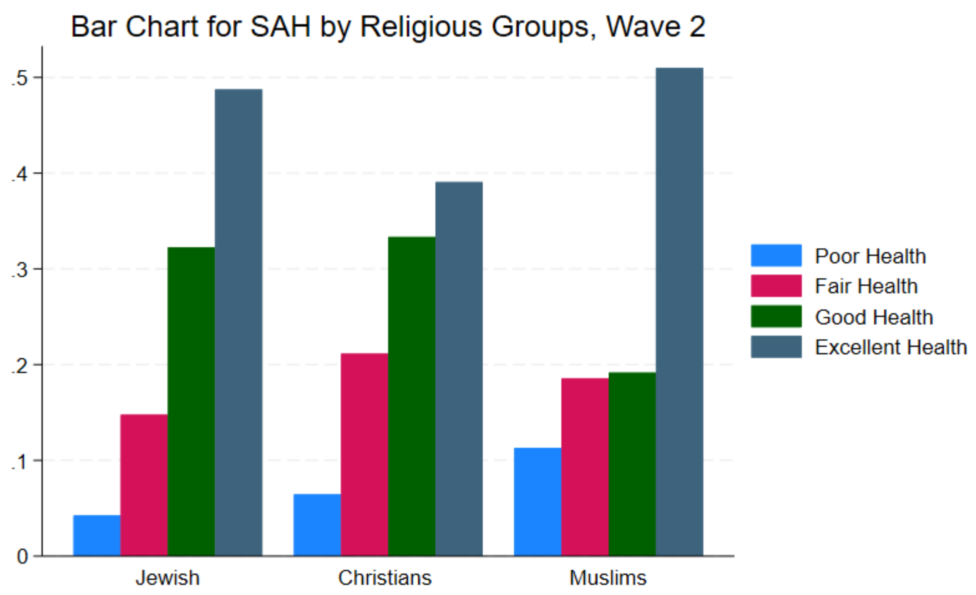

with a one-unit increase in lagged log-income and the presence of an academic degree. The results indicate that a one-unit increase in lagged log-income reduces the probability of reporting poor health by approximately 0.01 percentage points, fair health by 0.06 percentage points and good health by 0.10 percentage points, while increasing the probability of reporting

Table 2.2 Absolute and Relative Income- Marginal Effects

	Dependent Variable: *SAH*			
	Poor Health	Fair Health	Good Health	Excellent Health
Without Relative Deprivation Measure				
Loglincome	-0.0001	-0.0006	-0.0010	0.0017
	(0.0002)	(0.0011)	(0.0020)	(0.0034)
Academic Degree	-0.0053***	-0.0293***	-0.0519***	0.0866***
	(0.0006)	(0.0030)	(0.0051)	(0.0086)
With Relative Deprivation Measure				
Loglincome	-0.00002	-0.00012	-0.00021	0.00035
	(0.0002)	(0.0011)	(0.0021)	(0.0035)
Dep	0.00137***	0.00783***	0.01391***	-0.0231***
	(0.0003)	(0.0018)	(0.0032)	(0.0054)
Academic Degree	-0.00487***	-0.02775***	-0.04931***	0.08194***
	(0.0006)	(0.0030)	(0.0053)	(0.0088)

N= 20,509

Note: Standard errors are calculated using the delta method.
*** Significantly different from zero at the 1% level.
** Significantly different from zero at the 5% level.
* Significantly different from zero at the 10% level.

excellent health by around 0.17 percentage points. Although these effects are not statistically significant at conventional levels, they suggest a positive association between income and self-assessed health. Economically, this pattern implies that higher-income individuals are marginally more likely to rate their health more favourably, consistent with the broader literature linking material well-being and health outcomes. The small magnitude of the effects may reflect the limited within-person variation in income over time and the overall affluence of the sample, which constrains the potential for large health disparities based on income alone.

The significance of academic degrees persists, albeit with a slight reduction in magnitude. The presence of an academic degree continues to exert a positive effect on excellent health and a negative effect on poorer health categories. Similar trends for other socio-demographic variables are observed, consistent with the model excluding deprivation measures. While the dummy variable of being a female remains statistically significant, its significance increases to the 5% level when the deprivation measure is included, as opposed to the 10% level observed in the model without deprivation.

Proficiency in Hebrew, Arabic, and English languages emerges as a statistically significant predictor of health status. Individuals with proficiency in Arabic and English are more inclined to report poorer health outcomes compared to those proficient in other languages. In addition, other demographic variables, such as age cohorts, also demonstrate significant effects on health status. As expected, the effect of age on self-reported health status decreases with age, becoming smaller for older age groups. Furthermore, the effect of age on self-reported health status remained statistically significant for each age group. Gender differences are also evident in health outcomes. *Female* exhibits a negative association with better health with 10% significance levels.

When the relative deprivation measure is incorporated into the analysis (*Table 2.2*), notable changes emerge in the magnitude of the marginal effects associated with lagged log-income. The probability of reporting fair health decreases by approximately 0.012 percentage points (down from 0.06), while the likelihood of reporting good health decreases by 0.002 percentage points (down from 0.10) for a one-unit increase in lagged log-income. Comparable attenuation is observed for the excellent health category, suggesting that once relative income differences are accounted for, the direct effect of absolute income on self-assessed health becomes smaller.

The relative deprivation measure itself shows a clear and statistically significant pattern: higher deprivation scores are associated with a greater probability of reporting poor or fair health, and a lower probability of reporting excellent health. This implies that individuals who perceive themselves as worse off than others in their reference group experience measurable declines in self-rated health, even after controlling for their absolute income level.

Table 2.3 presents the binary outcome variable of good health and summarises the effects of variables of interest (income and the deprivation measure) on the outcome for the applied models: a linear probability model is estimated via ordinary least squares (OLS), probit with random- and fixed-effects using the Mundlak approach. It reports the marginal effects of probit models. Subsequently, the Hausman-Taylor estimator is included and the analysis concludes with the semiparametric approach, where income is assumed to have an unspecified relationship with good health status.

The results consistently indicate that lagged-log income positively affects good health. This effect diminishes and becomes statistically insignificant in Models 5–8. When the deprivation measure is included in the Hausman-Taylor estimator, the effect of income turns negative. However, this is observed only in Model 8, when including a relative deprivation.

The association between income deprivation measure and health is more pronounced, showing a negative and statistically significant relationship in all models. The effect is marginally highest when assuming a non-linear relationship between income and good health. However, the trend remains consistent compared to models with parametric assumptions.

Table 2.3 Robustness Check: Comparison of different models

	OLS (1)		OLS (2)		Probit RE ME(3)		Probit RE ME(4)	
	Coef.	Std. Err.	Coef.	Std. Err.	dy/dx	Std. Err.	dy/dx	Std. Err.
loglincome	0.0268***	0.0019	0.0185***	0.0023	0.1281***	0.00185	0.00665***	0.002
dep			-0.0023***	0.0003			-0.0031***	0.0004

	Probit Mundlak ME(5)		Probit Mundlak ME (6)		Hausman-Taylor (7)		Hausman-Taylor (8)		SP (9)	
	dy/dx	Std. Err.	dy/dx	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
loglincome	0.0006	0.0025	0.0002	0.0025	0.0002	0.0024	-0.0008	0.0024		
dep			-0.0024***	0.0005			-0.0017***	0.0004	-0.0023***	0.0037

2.5.3 Sub-samples of Religious Groups

Table 2.4 presents the findings of an analysis that examines the impact of absolute income on health outcomes across Jewish, Muslim, and Christian communities. Across all religious groups, income shows a noteworthy association with health status. Specifically, among Jewish individuals, higher income levels are associated with an improved health status. Muslims participants experience a negative income effect on health, where higher income levels are positively linked with poorer health outcomes, particularly evident in the Poor, Fair, and Good health categories. Christians exhibit a similar trend, with income demonstrating comparable results, albeit with statistical significance of 5% levels.

Among Jewish individuals, a one-unit increase in lagged log-income slightly reduces the probability of reporting poor, fair, and good health by 0.03, 0.14, and 0.14 percentage points, respectively, while increasing the likelihood of reporting excellent health by 0.32 percentage points. For Muslims, the same income increase raises the probability of reporting poor, fair, and good health by 0.03, 0.14, and 0.14 percentage points, respectively, but lowers the probability of reporting excellent health by 0.32 percentage points. The Christian sub-sample exhibits the largest income-related differences across health categories, although the estimates are imprecisely measured.

Having an academic degree displays a similar pattern across religious groups. Among Jewish individuals, holding a degree reduces the probability of reporting poor, fair, and good health by 0.4, 2.3, and 4.9 percentage points, respectively, while increasing the likelihood of reporting excellent health by 7.5 percentage points. Among Muslims, these effects are of

Table 2.4 Absolute Income, Religious Sub-samples - ME

	Dependent Variable: *SAH*			
	Poor	Fair	Good	Excellent
Jewish (N = 14,409)				
Loglincome	-0.00001	-0.00007	-0.00015	0.00023
	(0.0002)	(0.0012)	(0.0027)	(0.0042)
Academic Degree	-0.0038***	-0.0227***	-0.0488***	0.0754***
	(0.0006)	(0.0030)	(0.0063)	(0.0098)
Muslim (N = 2,293)				
Loglincome	0.0003	0.0014	0.0014	-0.0032
	(0.0012)	(0.0050)	(0.0053)	(0.0117)
Academic Degree	-0.0134***	-0.0525***	-0.0553***	0.1212***
	(0.0041)	(0.0141)	(0.0146)	(0.0320)
Christian (N = 629)				
Loglincome	0.0061**	0.0297**	0.0209**	-0.0567**
	(0.0028)	(0.0120)	(0.0086)	(0.0226)
Academic Degree	-0.0068	-0.0322	-0.0233	0.0634
	(0.0047)	(0.0216)	(0.0150)	(0.0407)

Note: Standard errors are calculated using the delta method.
*** Significantly different from zero at the 1% level.
** Significantly different from zero at the 5% level.
* Significantly different from zero at the 10% level.

greater magnitude, with decreases of 1.3, 5.2, and 5.5 percentage points in poor, fair, and good health, and an increase of 12.1 percentage points in excellent health. For Christians, the direction of the effects is similar but statistically insignificant. Economically, these results suggest that higher education is strongly associated with better self-assessed health, and that income–health gradients differ somewhat across religious groups, potentially reflecting variation in socio-economic context and reference-group perceptions.

Additional sociodemographic factors were analysed in religious sub-samples (*Table B.2*). Factors such as marriage, gender, and language exhibit varying significance among religious groups. For instance, marriage is statistically significant only for Muslims, while the effect of being female is negative for Jews but positive for Muslims and Christians. These findings highlight the varying impact of income, education, and sociodemographic factors on health across religious groups.

Table 2.5 extends the analysis by adding the relative deprivation measure, which significantly affects health outcomes. For Jewish and Muslim participants, deprivation worsens

Table 2.5 Absolute and Relative Income, Religious Sub-samples - ME

	Dependent Variable: *SAH*			
	Poor	Fair	Good	Excellent
Jewish (N = 14,209)				
Loglincome	0.00001	0.00004	0.00009	-0.00014
	(0.0002)	(0.0013)	(0.0027)	(0.0043)
Dep	0.001***	0.006***	0.013***	-0.020***
	(0.001)	(0.002)	(0.004)	(0.007)
Academic Degree	-0.003***	-0.021***	-0.045***	0.069***
	(0.001)	(0.003)	(0.006)	(0.010)
Muslim (N = 2,234)				
Loglincome	0.0004	0.0016	0.0018	-0.0039
	(0.0012)	(0.0050)	(0.0054)	(0.0118)
Dep	0.002	0.011	0.012	-0.027
	(0.002)	(0.011)	(0.011)	(0.025)
Academic Degree	-0.009**	-0.040***	-0.043***	0.092***
	(0.004)	(0.014)	(0.015)	(0.033)
Christian (N = 610)				
Loglincome	0.0067**	0.0333***	0.0242**	-0.0643***
	(0.0030)	(0.0126)	(0.0093)	(0.0237)
Dep	-0.003	-0.017	-0.013	0.0335
	(0.004)	(0.022)	(0.016)	(0.042)
Academic Degree	-0.006	-0.034	-0.025	0.066
	(0.004)	(0.021)	(0.015)	(0.041)

*** Significantly different from zero at the 1% level.
** Significantly different from zero at the 5% level.
* Significantly different from zero at the 10% level.

health outcomes, although the effect is statistically insignificant for Muslims. Among Christians, deprivation has the opposite effect, improving health in poorer categories.

The academic degree effect remains consistent between all religious groups when considering relative deprivation. However, marginal effects vary among religious groups, with Christians showing the strongest link between deprivation and health outcomes (*Table B.2*). The inclusion of deprivation consistently reveals similar patterns across sociodemographic factors, but deprivation and income have a more pronounced impact on the health of Christians compared to other groups.

2.6 Discussion

2.6.1 The relationship between health and absolute income

The extensive literature supports the absolute income hypothesis affecting health (Grossman, 1972; Jones and Wildman, 2008; Preston, 1975; Subramanian and Kawachi, 2004), confirmed by this study using Israeli data, showing a positive correlation between income and health. Despite complexities and inconsistencies in previous research (Frijters et al., 2005) (e.g. measurement errors and cross-sectional data issues), most studies, including this one, find a positive link between income and self-reported health (Benzeval and Judge, 2001; Gravelle and Sutton, 2009). The results in *Table 2.2* reaffirm this, with robustness checks validating the findings.

Other econometics applications such as the Hausman and Taylor (Hausman and Taylor, 1981) and semiparametric approach using a local polynomial fit (Robinson, 1988) revealed consistent results with those of the parametric specification, providing validation for the observed effects and supporting the presence of a non-linear relationship between income and health.

2.6.2 The relationship between health and relative income

The effect of relative income on health is debated, with studies showing mixed results (Clark and Oswald, 1996; Mangyo and Park, 2011; McBride, 2001; Senik, 2008). For example, Clark and Oswald (1996) and Senik (2008) found varying impacts of reference income on satisfaction across European countries. Our study found that absolute income's effect on health is more pronounced without the relative deprivation measure, but this was not statistically significant. Greater deprivation rank consistently led to worse health across models, especially in the fixed-effects and semiparametric analyses. Unlike the impact of absolute income on health, the effect of relative income was statistically significant in our primary model as well as across all other models, including those in the robustness analysis.

The findings align with recent empirical research, including Cai et al. (2022), which examined the effects of absolute and relative income on health outcomes in China. Their study demonstrated that both income measures exert significant and independent influences on individual health, with relative income playing a particularly important role in shaping self-reported health in both rural and urban populations. Although some variation was observed between sub-groups, the effects were especially pronounced among the rural population across all income inequality indicators applied. These results highlight the importance of incorporating relative income alongside absolute income in analyses of health determinants. Moreover, the study underscores the value of conducting subgroup analyses, as the strength and direction of income-health relationships may vary considerably across different demographic or socio-economic contexts.

We estimate a range of alternative model specifications to examine the sensitivity of the results to different assumptions and functional forms. Rather than relying on a single preferred specification a priori, our primary approach is to interpret the findings jointly across models and to assess whether the estimated effects are consistent in sign, magnitude, and statistical significance. In practice, the OLS specification serves as a natural benchmark due to its transparency and ease of interpretation, while the semi-parametric model relaxes functional form assumptions on income and provides a useful robustness check. The close alignment of results across these specifications increases confidence that the main conclusions are not driven by restrictive distributional or linearity assumptions, but instead reflect a stable underlying relationship.

The inclusion of both absolute and relative income measures raises potential concerns about multicollinearity, as individuals with higher absolute income are typically also positioned higher in the income distribution. In the present sample, this correlation is evident, and its primary manifestation is a reduction in the precision of the estimated coefficient on absolute income when relative deprivation is added to the model. As shown in Table 2.2, the point estimates for absolute income remain stable across specifications, but their standard errors increase modestly and the coefficients lose statistical significance once relative deprivation is included. This pattern suggests that while the correlation affects precision, it does not appear sufficiently severe to destabilise estimation or generate implausible coefficient magnitudes. Importantly, the relative deprivation measure remains precisely estimated and statistically significant across all health categories, indicating that its association with self-assessed health is not driven by collinearity with absolute income. Accordingly, the results are interpreted with an emphasis on robustness across specifications rather than reliance on a single coefficient estimate.

The sub-sample analysis revealed significant differences across religious groups, with income effects most pronounced for Christians *(Table 2.4* and *Table 2.5*). Our findings similarly demonstrate that subgroup-specific patterns, such as those observed across religious groups, can yield indispensable insights that would be obscured in aggregate-level models. This further reinforces the need for nuanced approaches when examining income-related health inequalities.

2.6.3 The relationship between income, health, and religious groups

The findings from the sub-sample analysis examining the relationship between absolute income and self-reported health status across three major religious groups in Israel: Jewish, Muslim, and Christian individuals. The results reveal that income is an important determinant of health across all groups, but the direction and nature of this relationship vary substantially, suggesting that cultural, behavioural, or structural factors may moderate the income-health gradient.

Among Jewish respondents, higher income is consistently associated with better self-reported health, a pattern that aligns with standard economic theories of health production (Grossman, 1972) and much of the existing literature on absolute income effects (Deaton, 2001; Subramanian and Kawachi, 2004). This positive association is expected in contexts where greater financial resources facilitate access to healthcare, healthier food, and lower exposure to psychosocial stressors.

Conversely, for Muslim individuals, the results suggest a counterintuitive pattern: higher income levels are associated with poorer self-reported health outcomes. This finding may reflect cultural or social dynamics that mediate the interpretation of health status or the use of healthcare resources. For example, it is possible that in some contexts income does not translate directly into improved health behaviours or access to care, or that upward economic mobility comes with increased stress or social obligations that negatively impact health (Clark and Oswald, 1996; Senik, 2008). A similar, albeit less pronounced, pattern is observed among Christian respondents, who also exhibit a positive association between income and lower health status in some categories and statistically significant[1].

The results also contribute to the broader literature on relative deprivation. When a measure of relative income is introduced, the pattern of findings shifts: among Jewish individuals, increased levels of deprivation relative to others within the same religious group are associated with a statistically significant decline in self-reported health. A similar relationship is observed among Muslim respondents. These findings are consistent with expectations of the relative income hypothesis, which posits that individuals evaluate their wellbeing not only in absolute terms but also through social comparisons within their reference groups (Jones and Wildman, 2008; Marmot et al., 1984; Wilkinson, 1996).

These differences show the importance of examining the role of religious and cultural identity when evaluating the income-health relationship. Religious affiliation may shape social comparisons, health beliefs, dietary practices, and even stigma around reporting health

[1] At 5% significance levels.

problems, all of which can influence the way individuals perceive and report their health. For instance, in Judaism, the consumption of wine and participation in ritual meals are common and sometimes positively associated with social cohesion and wellbeing. Conversely, in some Islamic contexts, higher income may be associated with more sedentary lifestyles or social expectations that unintentionally increase health risks.

These findings suggest that the income-health relationship is not uniform across cultural or religious lines. Public health interventions and income-based policies must therefore be sensitive to the social and cultural contexts in which individuals are embedded. A one-size-fits-all approach to improving health through income redistribution may overlook these group-specific dynamics and fail to address the underlying sources of health inequality.

2.7 Conclusion

This study supports both the absolute and relative income hypotheses in their relationship to health. Consistent with previous research (Grossman, 1972; Jones and Wildman, 2008; Subramanian and Kawachi, 2004), higher income correlates positively with better self-reported health, as shown using a rare Israeli dataset, despite complexities such as measurement errors (Frijters et al., 2005).

Across all specifications, log-lagged income exhibited a positive marginal effect on self-reported health; however, statistical significance was achieved only in the OLS models and in the random-effects probit specifications both with and without the deprivation measure. Once individual-level fixed-effects were incorporated using the Mundlak approach in the probit model, the coefficient on income lost statistical significance, indicating that unobserved heterogeneity, such as stable personality traits or long-term health behaviour, can confound the estimated effect in simpler models. This finding aligns with the argument that income-health gradients may be overstated in models that fail to account for individual fixed-effects.

In contrast, the relative income hypothesis, captured through a rank in the income distribution, yielded more robust evidence. The deprivation measure was statistically significant across all specifications, indicating that individuals' relative position in the income distribution has a meaningful effect on their perceived health. Notably, the magnitude and significance of the deprivation effect remains consistent across both OLS and semiparametric estimations, suggesting robustness to model specification and indicating that nonlinearity does not significantly influence the relationship in this case.

The religious subgroup analysis revealed considerable heterogeneity. When only absolute income was included, Muslims and Christians exhibited a similar positive association with health. However, after introducing the relative deprivation measure, the patterns shifted. Jews and Muslims displayed more comparable trends in this model, both showing stronger sensitivity to relative deprivation, with Muslims experiencing the largest negative effects. These differences may reflect variation in social comparison norms, expectations about material well-being, or community structures. By comparing individuals with others within their religious group as well as in the overall population, this study provides new insights into the contextual salience of reference groups, contributing to the broader literature on relative income and well-being (Clark and Oswald, 1996; Mangyo and Park, 2011; Senik, 2008).

Despite the richness of the dataset, several limitations must be acknowledged. First, the use of self-reported health introduces potential reporting bias, particularly across cultural or religious groups with differing perceptions or stigma surrounding health. Second, the sample of Christian participants is relatively small, limiting the precision of subgroup estimates for that group. Third, the panel includes only a limited number of waves, which constrains the analysis of long-term dynamics and the potential lagged effects of income or deprivation on health.

In general, the findings provide strong support for the role of relative income in shaping health outcomes. The stronger and more consistent effects observed for relative deprivation underscore the need to go beyond income redistribution alone and address the broader social and psychological dimensions of inequality. Policymakers should take into account not only material deprivation, but also the perceived fairness and visibility of income differences when crafting health and social policies. Interventions targeting vulnerable communities, especially where social comparisons are most salient, may be particularly effective. These results reinforce the growing consensus that economic policies cannot be disentangled from public health goals, particularly in contexts characterised by high inequality or cultural diversity.

Chapter 3

The Effect of Obesity on Employment: Using Genetic Variants from UK Biobank as an Instrumental Variable

AARON KATZ

Abstract

This study explores the causal impact of obesity on employment using data from the UK Biobank, a comprehensive dataset combining genetic, health, and socioeconomic information. We apply a Mendelian Randomisation (MR) approach, using a genetic risk score (GRS) for a higher body mass index (BMI) as an instrumental variable. This method leverages genetic predispositions determined at conception, enabling us to isolate the exogenous variation in obesity while addressing key endogeneity concerns, such as omitted variable bias and reverse causation. To account for the heterogeneity in obesity's impact across diverse demographic and socioeconomic factors, we apply a combination of IV techniques and the person-centered treatment effects (PeT) framework. This methodological approach extends beyond traditional average treatment effect estimations, uncovering nuanced individual-specific impacts that reflect the complex relationship between obesity and employment probabilities. Our findings reveal a significant and negative causal effect of obesity on employment across all estimation methods. Nonlinear approaches, including marginal treatment and PeT effects, yield the same substantive results with differences in marginal effects. Gender-specific subsample analyses indicate that the employment penalty associated with obesity is more pronounced among females in the linear models, whereas the PeT effects approach reveals the opposite pattern. Overall, these results highlight the wider economic burden of obesity, particularly its

adverse impact on employment.

Keywords: Obesity, employment, Mendelian randomisation, genetic risk score, BMI, UK Biobank, person-centered treatment effect

JEL Codes: I01, I12, J01, J7, J21, C26

3.1 Introduction

The rising prevalence of obesity has sparked significant concern over its widespread economic and social consequences. One critical issue is its impact on employment. A growing body of literature has documented a negative association between obesity and employment, including lower wages, reduced employment probabilities, and increased absenteeism (Baum and Ford, 2004; Busetta et al., 2020; Cawley, 2004; Li et al., 2022; Morris, 2007). These adverse effects stem from multiple mechanisms, such as weight-based discrimination, early labour market withdrawal, and health-related productivity losses and cost.

Obesity-related illnesses result in an estimated four additional lost workdays per employee with obesity each year, placing a significant economic burden on employers and public health services (Destri et al., 2022). Evidence shows that unemployment rates are significantly higher among individuals with obesity(Bajorek and Bevan, 2019). In response, the UK government has considered offering GLP-1 receptor agonist treatments, such as Ozempic, to unemployed people with obesity as part of a strategy to support their return to the labour market. This reflects the high cost of obesity to companies, including an estimated £88 billion in medical expenses, £87 billion in presenteeism, and £63 billion in absenteeism (Dall et al., 2024; Diabetes UK, 2024). For a hypothetical company with 10,000 employees, the annual cost of obesity has been estimated to be £13 million, with an average of around £5,000 per affected employee (Dall et al., 2024). Despite these well-documented associations and costs, establishing a causal relationship remains difficult due to potential endogeneity, arising from omitted variable bias and reverse causality (Böckerman et al., 2019; Cawley, 2004; Norton and Han, 2008). Understanding these causal pathways is essential for designing effective policies to mitigate the economic repercussions of rising obesity rates.

Our study contributes to the existing literature through three main channels. First, to illustrate endogeneity concerns, we examine the effect of obesity[1] on the probability of being in paid employment using a range of estimators. One approach involves including a Genetic Risk Score (GRS) for obesity (measured by BMI[2]), as an IV to estimate the causal effect of obesity on employment, following the Mendelian randomisation framework (Böckerman et al., 2019; Sanderson et al., 2022), using the UK Biobank data. Genetic variants associated with BMI are fixed at conception and remain unaffected by socioeconomic or behavioural factors. This makes the GRS a credible instrument for isolating exogenous variation in obesity (Lawlor et al., 2008a), strengthening the causal interpretation of our findings.

[1]Constructed from BMI values equal to or greater than 30.

[2]Calculated using the formula: $BMI = \frac{\text{weight (kg)}}{\text{height (m)}^2}$.

The attempts in the existing literature to identify causal effects are largely based on conventional IV methods, which can yield inaccurate estimates when treatment effects vary over unobserved confounders across individuals. Many studies have used the weight of a biological relative as an IV for an individual's weight (e.g., Cawley, 2004; Lindeboom et al., 2010). A key limitation of using a biological relative's BMI as an instrument is the potential for unobserved confounding factors that can influence both an individual's BMI and their relative's. For example, shared environmental and socioeconomic conditions, such as diet, education level, and income, can influence both individuals, violating the exclusion restriction and introducing bias in estimation. To address this issue, using a GRS for high BMI as an instrument offers a more robust approach. Since GRS is fixed at conception and randomly allocated[3], it is less likely to be correlated with environmental or behavioural confounders that affect labour market outcomes (Davey Smith and Ebrahim, 2003). This makes it a valid instrument under the assumptions of Mendelian Randomisation, where genetic variants are used as proxies for modifiable exposures (in this case, obesity).

A second contribution to the existing literature is that, given that the effect of obesity may vary depending on both observed and unobserved factors, IV methods (even when using a strong and valid instrument) may produce biased estimates of the average treatment effect (ATE) or the conditional average treatment effect (CATE), due to essential heterogeneity (Basu, 2014; Basu et al., 2018). Essential heterogeneity refers to the idea that individuals self-select into treatment (or exposure) based on anticipated gains or losses, which are influenced by unobserved characteristics. In this context, individuals may respond differently to increases in BMI based on factors such as genetics, health status, or socioeconomic background, making the treatment effect inherently non-random across the population. As a result, IV estimates may reflect a local average treatment effect (LATE) that does not generalise to the broader population.

The ATE represents the average effect of a treatment across the entire population, whereas the CATE captures the effect of treatment for specific subgroups, conditioned on observed characteristics. When treatment effects are heterogeneous, IV methods may fail to identify the ATE because they primarily estimate a LATE, which applies only to individuals whose treatment status is influenced by the instrument. This limitation arises because unobserved factors, such as an individual's time preference (discount rate) or motivation to work, may moderate the effect of obesity on employment via the instrument (i.e., compliers[4]), leading to variation in treatment effects across individuals. Furthermore, these unobserved confounders

[3] Under Mendel's law.

[4] Individuals whose treatment status is influenced by the instrument.

not only induce selection bias (Newhouse and McClellan, 1998), but also contribute to treatment effect heterogeneity, as individuals typically make treatment decisions based on private information about their expected benefits (Heckman and Vytlacil, 1999). To address this issue, a local instrumental variable (LIV) approach is applied.

LIV and LATE approaches both aim to estimate causal effects in the presence of endogeneity but differ in their assumptions and scope. LATE estimates the causal effect of a treatment for compliers. Provides a local effect rather than a population-wide estimate. In contrast, LIV extends this by allowing for a more flexible estimation of heterogeneous treatment effects across different values of the instrument. LIV does not assume a binary instrument and instead estimates the full distribution of treatment effects, making it more informative in settings with continuous or multi-valued instruments.

Building on the continuous nature of our IV, we include LIV methods to estimate marginal treatment effects (MTEs), which offer a nuanced approach to capture treatment effect in the presence of essential heterogeneity. MTEs identify effects for individuals whose observed characteristics (e.g., BMI levels and genetic risk scores) interact with unobserved confounders, including factors such as discount rates and motivation, making them indifferent to have obesity or not. By using LIV methods, this approach accounts for endogeneity and provides deeper insights into how unobserved factors influence employment decisions via obesity. In other words, MTEs account for the fact that the impact of obesity on employment may not be uniform but depends on underlying personal traits that are not directly observed in the data. Furthermore, we include a person-centered treatment (PeT) effects approach, a weighted extension of MTEs (Basu, 2014; Basu et al., 2018), to estimate the impact of obesity on employment.

PeT effects provide a more comprehensive understanding of treatment-effect heterogeneity, capturing a greater degree of individual variability. In addition, it offers a more accurate measure of self-selection, improving the prediction of the true effect of treatment (Basu, 2014). The traditional analytical methods, which typically focus on average treatment effects or local treatment effects, may not accurately reflect the complexities of individual responses, potentially leading to misleading conclusions. Another key advantage of PeT effects is their flexibility; they allow for the direct computation of mean treatment-effect by sub-groups, such as the effect on the treated and other sociodemographic factors, making them a valuable tool for causal inference.

Although much of the existing literature on treatment effects has traditionally focused on estimating parameters that inform decisions at the population or policy level, the PeT effects approach may overlook the value of treatment effects at the individual level, particularly in the presence of treatment effect heterogeneity (Basu, 2014; Basu et al., 2018; Heckman and Vytlacil, 1999). For example, individuals may respond differently to changes in obesity status with respect to being in employment, depending on unobserved preferences or characteristics that influence their likelihood of being a person with obesity. As such, the socially optimal intervention (e.g., a universal health subsidy or employment support for individuals with obesity) may not align with the choices individuals would make if fully informed about their potential employment outcomes. Moreover, if treatment-effect information (in this case, the causal impact of obesity on employment status) can guide individuals toward more informed choices, such as engaging in health-enhancing behaviours, it could lead to a form of positive self-selection (reducing the probability of being an individual with obesity) and improves both individual welfare and overall policy effectiveness.

Finally, the potential for measurement error in obesity, as captured by BMI, raises concerns about the accuracy of this indicator. To address this issue, we re-estimate our models using an alternative measure, waist-to-height ratio, to assess the robustness of the results under the same identification strategy and estimation methods. Notably, this approach has not been previously implemented in the literature in conjunction with the methods applied here. Furthermore, since previous studies examining the relationship between obesity and employment have identified heterogeneous effects by gender (Böckerman et al., 2019; Caliendo and Gehrsitz, 2016; Cawley, 2004; Morris, 2007; Tyrrell et al., 2016), largely attributed to gender-based disparities in weight penalties (e.g., wages), this study investigates possible gender differences in the impact of obesity on the probability of being in paid employment. The primary analysis is conducted on gender-specific subsamples to determine whether the effect of obesity varies between men and women.

Our findings provide strong evidence of a substantial causal effect of obesity on the likelihood of being in a paid employment, consistently across all estimation methods. In particular, the effect becomes more pronounced when accounting for endogeneity, highlighting the importance of addressing potential bias in observational data. The effect follows a similar pattern when obesity is measured by waist-to-height ratio. In addition, we discover meaningful heterogeneity in the impact of obesity by gender at the individual level. Analysis of PeT effects reveals that the relationship between obesity and employment is marginally greater compared to the other estimation methods.

The findings of PeT effects reveal how the impact of obesity varies across individuals with different levels of unobserved resistance or susceptibility to treatment. The PeT estimates show that the negative effect of obesity on employment is not uniform across the population, it tends to be larger for individuals who are more likely to be with obesity based on their observed and unobserved characteristics (i.e., those with higher propensity scores). This suggests that the employment penalty associated with obesity is concentrated among individuals whose characteristics, both observed and unobserved, make them more likely to be with obesity.

Furthermore, we find heterogeneity in other sociodemographic factors, such as educational attainment. Specifically, individuals with lower levels of education who are also having obesity experience a greater employment penalty, with the interaction between obesity and lack of a degree associated with a significantly lower probability of employment. In contrast, having a degree appears to mitigate some of the adverse effects of obesity on employment outcomes, with the interaction term showing a positive association with being employed. This suggests that the adverse labour market effects of obesity are exacerbated among socioeconomically disadvantaged groups, further highlighting the interaction between health, human capital, deprivation status, gender, and employment outcomes. These findings show the value of methods like PeT that account for unobserved heterogeneity, go beyond the average treatment effect, and offer a more policy-relevant understanding of distributional impacts.

This paper is structured as follows. Section 3.2 presents a review of previous literature on the association between obesity and employment. Section 3.3 outlines the data source used in the presented study. In Section 3.4, the methods are introduced, including a review of the IV assumptions and the MTE and PeT effects, while Section 3.5 presents the results of the main model. Section 3.6 discusses the results with a comparison with the previous literature and concludes the paper.

3.2 Literature Review

The relationship between obesity and employment has been extensively studied, yet establishing causality remains a challenge. Lundborg et al. (2006) investigated the relationship between obesity and employment probability using data from the Survey of Health, Ageing, and Retirement (SHARE) in Europe, which tracks individuals over 50 years of age from ten European countries. Their findings indicate that obesity significantly reduces the likelihood of being employed. Cawley (2004) using US data found that obesity is negatively associated with wages and employment probabilities, particularly for women.

Baum and Ford (2004) extend the above findings by conducting a longitudinal analysis of the wage effects of obesity. Their study confirms that the wage penalty persists over time and is particularly pronounced for individuals who remain with obesity throughout the study period. Morris (2007) showed that weight gain is associated with subsequent declines in employment, although reverse causality could not be completely ruled out. Furthermore, Dall et al. (2024) and Han et al. (2011) investigate both the direct and indirect effects of body weight on adult wages, finding that obesity not only directly reduces wages, but also indirectly affects them through lower educational attainment and poorer health.

Despite robust evidence on the association between obesity and adverse employment, establishing a causal relationship has proven challenging due to potential endogeneity issues. These issues arise from omitted variable bias and reverse causality. On the one hand, an individual's employment status could significantly impact their dietary choices and physical activity patterns, consequently influencing their BMI status (Cawley, 2004; Kim and Leigh, 2010). Conversely, an individual's health condition, often reflected in their BMI, may notably affect their ability to secure employment (Goulão et al., 2024). Furthermore, endogeneity could also arise from unobserved factors such as individual preferences or other unobservable factors that simultaneously influence both employment and BMI, adding another layer of complexity to the analysis. To address these limitations, researchers have applied various methodological approaches.

Recent studies have leveraged genetic instruments to provide more robust causal estimates. Mendelian Randomisation (MR) is a technique that uses genetic variants associated with risk factors (such as BMI) as instruments to infer causality (Böckerman et al., 2019; Sanderson et al., 2022; Tyrrell et al., 2016). The use of the Genetic Risk Score (GRS) as instrumental variables represents a promising avenue to isolate exogenous variation in obesity. Genetic variants associated with BMI are determined at conception and therefore are not influenced by confounding socioeconomic or behavioural factors, which makes them robust instruments for estimating the causal effect of obesity on employment status.

Von Hinke et al. (2016) outline three key assumptions that must be satisfied when using genetic variants as instrumental variables: relevance (the genetic variant must be robustly associated with the exposure of interest, such as obesity), independence (the instrument must be independent of unobserved confounders that affect the outcome), and exclusion restriction (the instrument affects the outcome only through its effect on the exposure). They demonstrate how these assumptions can be evaluated in practice using two empirical applications, one examining the effect of a child's fat mass on academic performance and the other on blood pressure. The latter provides convincing evidence that a higher fat mass causally increases blood pressure, supporting the validity of the instrumental variable strategy. Building on this framework, several studies have applied similar approaches to estimate the causal effect of obesity on labour market outcomes (Böckerman et al., 2019; Campbell et al., 2021; Tyrrell et al., 2016; Von Hinke et al., 2016).

Campbell et al. (2021) use MR techniques to estimate the impact of increased BMI on employment status in the UK, finding robust evidence that higher genetically-predicted BMI reduces the probability of being employed, particularly among women. Similarly, Böckerman et al. (2019) apply the MR method in a Finnish context and report that a higher weight significantly reduces earnings and employment probabilities, providing strong support for a causal link between adiposity and adverse labour market outcomes. Tyrrell et al. (2016), while primarily focused on maternal obesity-related traits and birth weight, contribute valuable evidence on the biological basis of obesity through genetic associations, offering important validation for the instruments used in MR analyses. Collectively, these studies highlight the importance of addressing endogeneity in the obesity-employment relationship and highlight the utility of MR approaches for causal inference in health economics.

Using GRS as an IV offers two distinct advantages: Firstly, it demonstrates greater efficacy in explaining variation in BMI compared to individual Single Nucleotide Polymorphisms (SNPs), which are the most common type of genetic variation among people, representing differences at a single position in DNA[5]. These variations can influence traits such as body weight. The genetic instrument used here captures more of the variability in an individual's BMI attributable to genetic factors, based on the findings from a large-scale meta-analysis involving nearly 340,000 individuals conducted by Locke et al. (2015). Secondly, it potentially exhibits higher validity by mitigating the risk of bias due to alternative biological pathways (pleiotropy) inherent in individual SNPs. This is attributed to the GRS representing a count of alleles linked with higher BMI, using appropriate weights depending on the strength of the specific alleles to predict obesity (as determined by previous literature and Genome-Wide Association Study (GWAS)). This is in contrast to specific alleles of particular SNPs, which encapsulates the cumulative impact of multiple genetic variants on BMI (Palmer et al., 2012; Smith, 2011). Speliotes et al. (2010) clarify that the SNPs included in the GRS for obesity are those known to influence weight through mechanisms directly related to BMI, such as appetite and insulin secretion or response (with $r^2 = 0.83$).

Recent studies leveraging MR techniques have advanced the field by addressing endogeneity concerns using genetic instrumental variables, such as genetic risk scores for higher levels of BMI (Böckerman et al., 2019; Caliendo and Gehrsitz, 2016; Campbell et al., 2021; Sanderson et al., 2022; Tyrrell et al., 2016). This method exploits the strong genetic influence

[5]Deoxyribonucleic acid (DNA) is the molecule that carries genetic information for the development and functioning of an organism.

on body weight, with research estimating that genetic factors account for approximately 40–70% of BMI variation (Locke et al., 2015). Some studies suggest a higher correlation, exceeding 70% (Seral-Cortes et al., 2021; Speliotes et al., 2010).

Using GRS as an IV for causal inference may yield biased estimates of ATE and CATE due to the presence of essential heterogeneity, implying that treatment effects vary across unobserved confounders (Basu et al., 2018). Traditional IV methods often assume a homogeneous treatment effect across the population, potentially overlooking this heterogeneity. To address this limitation, advancements in econometric techniques, particularly the person-centred treatment effects (PeT) approach (Basu, 2014; Basu et al., 2018), offer a more flexible framework for capturing variation in treatment effects. Originally conceptualised by Heckman and Vytlacil (1999) within the Local Instrumental Variable (LIV) framework, PeT effects extend this approach by allowing a more personalised estimation of treatment effects, taking into account the nuanced ways in which treatment effects differ across individuals.

PeT effects approach includes an IV framework, conditioned on an individual's observed characteristics, and averaged over the potential conditional distribution of unobserved characteristics that influence their observed treatment. For instance, two individuals with similar genetic risk scores, levels of obesity, and observable factors may still experience different impacts on their employment outcomes due to varying unobservable factors. This methodology offers a better understanding of how genetic predispositions to higher BMI influence employment probabilities, taking into account the differential impacts of various sociodemographic factors (Basu, 2014).

PeT effects were previously used in the study of Basu et al. (2018), which examines the heterogeneous impact of educational reforms on long-term health outcomes. The authors leverage variation in exposure to selective and non-selective schooling systems in England and Wales to estimate the impact on depression, smoking, and other health-related behaviours. They find that the transition from a selective early-tracking system to a non-selective one significantly increased depression and cigarette smoking for a subset of individuals, with these effects persisting over time; notably, cognitive abilities did not moderate the effects, but those with lower non-cognitive skills were most adversely affected.

A recent study by Moler-Zapata et al. (2023) examined the performance of local instrumental variable approaches compared to the traditional 2SLS method, particularly in relation to the strength of the IV and variations in sample size, using a simulation study. The findings revealed that the LIV methods produced estimates with low bias across all samples, regardless of sample size, as long as the IV was strong. In contrast to 2SLS, LIV provided estimates for ATE and CATE with lower bias and root mean squared error.

Building on previous research that highlights significant differences in the impact of obesity on employment between genders (Böckerman et al., 2019; Busetta et al., 2020; Caliendo and Gehrsitz, 2016; Cawley, 2004; Goulão et al., 2024; Morris, 2007), our analysis adopts a sex-disaggregated approach to provide more precise and policy-relevant insights. A growing body of literature suggests that obesity affects employment differently for men and women, potentially due to factors such as gender-specific discrimination in the labour market, societal beauty norms, and occupational sorting (Kühn and Wolbring, 2024).

3.3 Data

The UK Biobank is a large-scale biomedical database that aims to improve the prevention, diagnosis, and treatment of a wide range of illnesses. Initially, individuals aged 38 to 73 were invited to participate in the study, providing detailed information about their health, lifestyle, and environment through questionnaires, interviews, and physical measurements. In addition, biological samples such as blood, urine, and saliva were collected for genetic analysis.

The research cohort comprised individuals who chose to volunteer for the UK Biobank, an ongoing study that includes approximately 500,000 adults living in England, Scotland, and Wales. These assessments included various measurements, such as measured height[6] and weight, enabling the calculation of BMI in kilograms per square meter. The mean BMI within our dataset of employed individuals is around 27.4 (*Figure 3.1*). For analysis, we construct a binary variable for obesity, defined as BMI above 30. This threshold is commonly used in published guidelines (NHS, 2023; NICE, 2014).

In defining an estimation sample, certain participants were excluded based on specific criteria, as can be seen in Figure 3.2. First, participants who were over the state retirement age at the time of initial assessment (60 for females and 65 for males during the first wave to the assessment centre) were excluded, as receipt of a state pension is an incentive to leave the labour market. Second, any participant who lacked data for employment status was not included in the analysis. Third, we include only participants who reported their ethnic background as 'White British'. This is because the genetic risk score for obesity in the UK Biobank is calculated based on a principal component analysis of genotypes, which requires a similar genetic ancestry for accurate results. This implies that the GRS for minority ethnic groups in the UK is likely to be unreliable. 59,855 individuals who reported not to be "White British" were dropped.

[6]Nurse measured by Seca 202 device, which is used primarily in hospitals, general practice and nursing facilities on wards.

Fig. 3.1 BMI among working individuals in the UK Biobank sample.

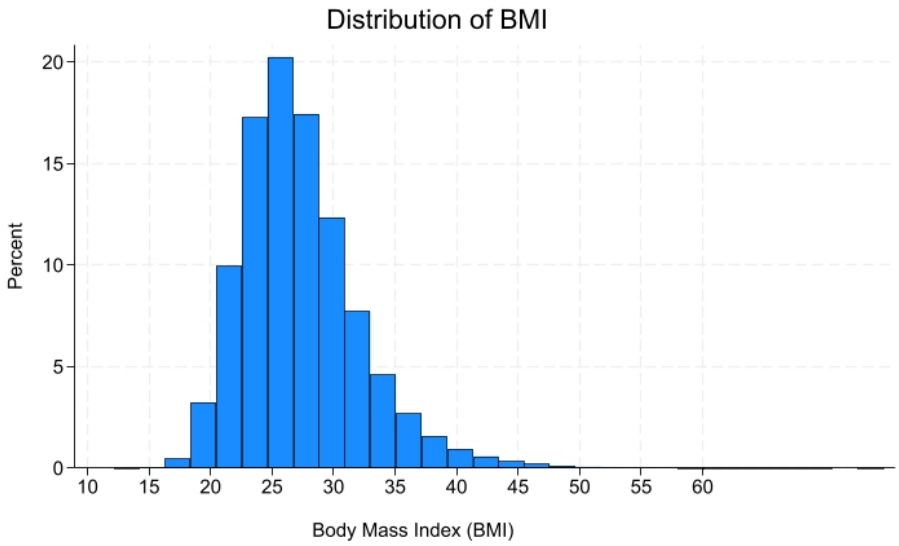

Genetic information is not available for all individuals in the dataset. After collapsing the data to the individual level, 16,229 individuals never provided valid genetic risk score information across all observed waves and are therefore excluded from analyses relying on genetic instruments. An additional 250 individuals were excluded from the sample after withdrawing consent for their data to be used in the UK Biobank, and were subsequently removed from the original dataset. Participants with missing BMI measurements or missing genetic risk scores related to BMI were also excluded. Among females, 276 were excluded for not reporting employment status, and 1,001 were excluded due to missing BMI data. Among males, 355 were excluded for not reporting employment status, and 1,254 were excluded for missing BMI measurements. Furthermore, 104,576 female participants and 43,343 male participants were excluded for being older than the state pension age on the date of their assessment centre visit. After applying all exclusion criteria, the final analytical sample consists of 284,270 individuals, comprising 130,387 females and 153,871 males.

Fig. 3.2 Sample construction and exclusion criteria.

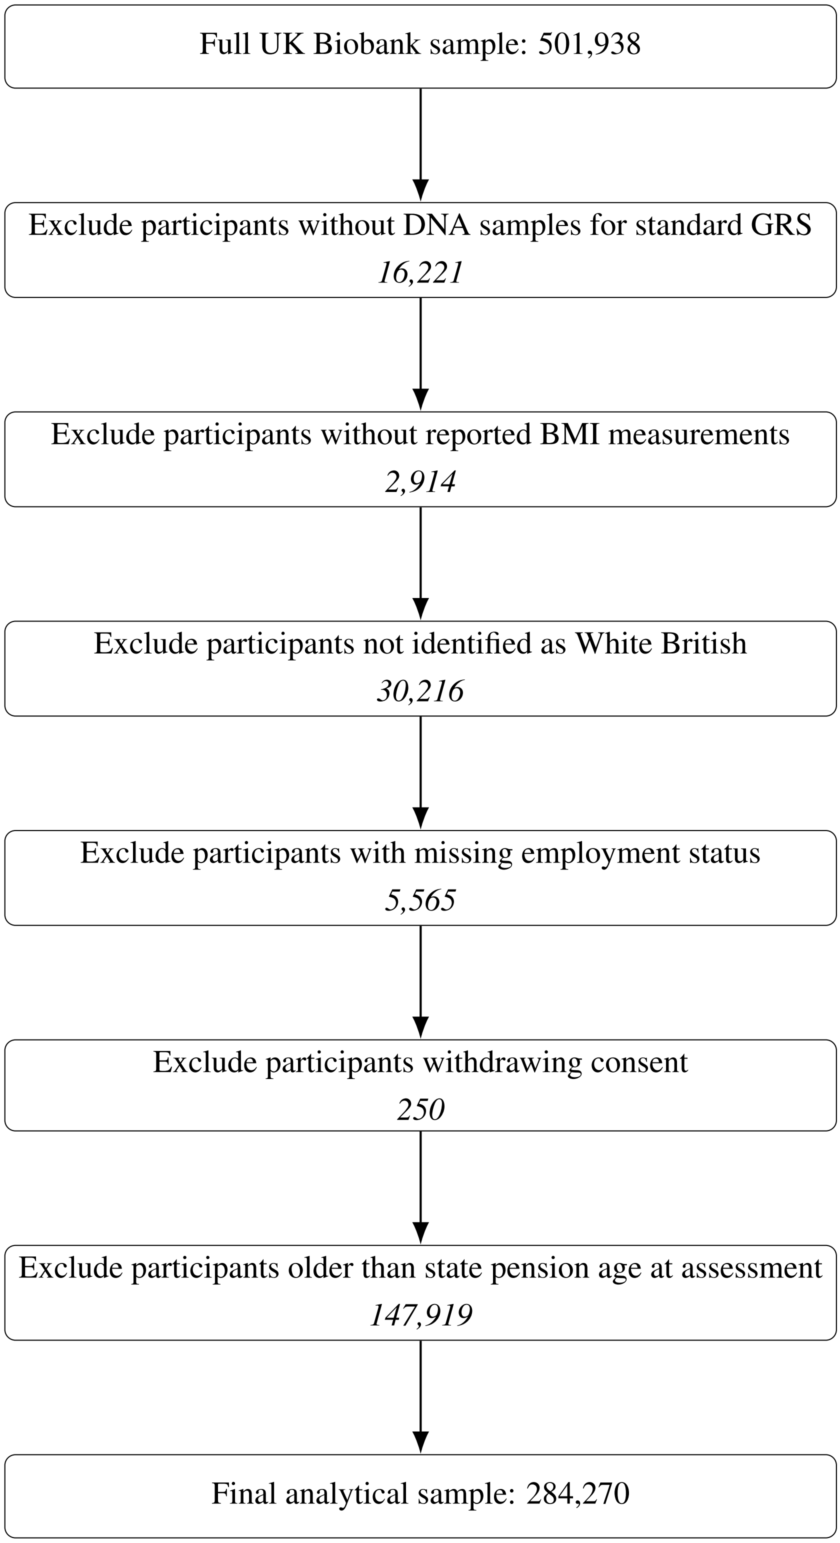

Employment status serves as our binary outcome variable. It is derived from a broader question in which participants report their current situation in the labour market, categorised into nine distinct groups. We construct a binary indicator equal to one for working-age individuals who report being employed or self-employed, and zero for all other statuses, including retirement, caring for relatives, or other non-employment categories.

We incorporate a Genetic Risk Score for obesity as an instrumental variable. The GRS is a metric used to assess an individual's risk of developing certain diseases based on 77 selected SNPs, those significantly associated with high BMI (with a stringent significance threshold of $p < 1.0 \times 10^{-8}$) (Thompson et al., 2022; Uffelmann et al., 2021). SNPs represent common genetic variations that arise from single nucleotide changes (A, T, C, or G) within the DNA sequence. These variations are distributed throughout the entire human genome, contributing to the remarkable diversity observed among individuals. Detecting SNPs involves techniques such as genotyping using microarrays or next-generation sequencing, as well as targeted Polymerase Chain Reaction (PCR) amplification followed by sequencing or restriction enzyme analysis (Brookes, 1999).

The construction of the GRS involves several key steps. The 77 SNPs used for constructing GRS for obesity serve as robust genetic markers distributed across the genome. Next, weights were assigned to each SNP based on its effect size on higher BMI levels. The weights were obtained from GWAS summary statistics and reflect effect sizes reported in previous studies, applied to genetic variants known to influence BMI. By aggregating the contributions of these SNPs, GRS was applied for each individual. Notably, the GRS directly represents an individual's relative risk compared to the general population due to population standardisation. The resulting standardised GRS values range from -4.709 (minimal genetic risk score) to 4.051 (maximal genetic risk score) in our sample of working-age individuals.

We control for a variety of indicators including gender, education level, geographical deprivation score, and age. Previous research investigating the relationship between weight and employment, as highlighted in Campbell et al. (2021), Böckerman et al. (2019), and Cawley (2004) revealed gender-based disparities. Cawley (2004) found that obesity is negatively associated with wages and employment probabilities, particularly for women. Given the observed heterogeneity in results by sex, this study conducts separate analyses for males and females to account for the different impacts of obesity across genders, as suggested in previous empirical work (Cawley, 2004; Conley et al., 2012).

3.4 Methods

This section outlines the methods used to evaluate the effect of obesity on employment status.

3.4.1 Baseline Model

The main model is specified as follows:

$$Y_i = \alpha + \beta_1 Obesity_i + \beta_2 X_{Oi} + \varepsilon_i, \tag{3.1}$$

where Y_i is a binary variable indicating whether an individual is in paid employment, $Obesity_i$ is a binary variable equal to one if the individual's BMI is 30 or higher and zero otherwise, X_{Oi} represents a vector of observed covariates, and ε_i is an idiosyncratic error term.

We begin by estimating Equation (3.1) using an Ordinary Least Squares (OLS), treating it as a linear probability model (LPM). This estimation captures the correlation between obesity and employment, but it may also be biased due to omitted variables and reverse causality.

Potential factors may contribute to a nonzero correlation between obesity, employment, and the error term. One such factor is the issue of reverse causality. Individuals experiencing unemployment may be prone to developing excess weight due to reduced income levels, potentially leading to the consumption of cheaper, less healthy food options (Drewnowski and Barratt-Fornell, 2004). Additionally, prolonged periods of unemployment can exacerbate feelings of depression, increasing the BMI levels and the risk of obesity, as indicated by studies such as Mossakowski (2009) and Blaine (2008).

Another factor is the problem of omitted variable bias. Certain unobserved individual characteristics, such as personality traits, time preferences and environment factors (e.g. supply of fast food), may be correlated with both employment status and obesity. For example, consider the case where an individual's intelligence or ability remains unobserved by the researcher. If ability is the only omitted variable, the error term in Equation (3.1) would consist of $\varepsilon_i = \theta A_i + \varepsilon_i$, where A represents ability and ε is an error term. It is plausible that individuals with lower (higher) ability levels are more (less) likely to be living with obesity (Kanazawa, 2013, 2014). This may be due to less informed dietary choices and a misunderstanding of the long-term consequences of excessive calorie consumption. Additionally, individuals with lower (greater) ability levels may also be less (more) likely to be employed. In such a scenario, the bias in the LPM estimator would likely be negative, as ability is expected to positively influence employment ($\theta > 0$), while obesity and ability are expected to be negatively related ($cov(\sigma_{Obesity,A}) < 0$). Consequently, LPM would lead to a negative relationship between obesity and employment biased away from zero.

To provide a further example, an employee may be characterised by an unobserved strong work orientation, exemplified by working after hours and dedicating more time to their work responsibilities. Given their high commitment to work, they may prioritise work over personal health, as extended working hours have been linked to an increased risk of obesity, as demonstrated by Virtanen et al. (2020). This unobserved behaviour will introduce a positive bias in the LPM estimator, resulting in the negative impact of obesity on employment being biased towards zero.

To address the endogeneity of obesity, we apply an IV approach using 2SLS estimation and a genetic risk score (GRS) for obesity as an instrumental variable:

$$Y_i = \beta_0 + \beta_1 \widehat{Obesity}_i + \beta_2 X_{Oi} + \varepsilon_i, \quad \text{where } Obesity_i \text{ is regressed on } X_{Oi} \text{ and } GRS_i \text{ via a first-stage.} \tag{3.2}$$

Here GRS_i represents the genetic risk score for obesity. The error term in the first-stage, u_i, captures unobserved factors influencing obesity, assumed orthogonal to X_{Oi} and GRS_i. ε_i is the error term capturing unobserved determinants of employment. Estimation by 2SLS ensures that standard errors are correctly adjusted for the two-stage estimation procedure.

3.4.2 Instrumental Variable Assumptions and Validity

To ensure the validity of the genetic risk score as an IV, several key assumptions must be satisfied.

Exclusion Restriction / Validity

The exclusion restriction requires that the instrument is uncorrelated with the error term ($cov(GRS, \varepsilon) = 0$), implying that GRS should only indirectly affect the outcome variable through the treatment variable. While no standardised test exists for this assumption, it holds if GRS influences employment solely via BMI levels rather than through any direct effect.

Previous studies provide evidence supporting this assumption when using genetic IV. Locke et al. (2015) and Speliotes et al. (2010) found that the SNPs associated with high BMI do not correlate with traits unrelated to obesity, such as cognitive ability, which could directly influence employment. Moreover, Böckerman et al. (2019) and Von Hinke et al.

(2016) found no evidence of a direct association between the SNPs used in the genetic risk score and employment status[7].

The SNPs selected to construct the GRS for BMI in the UK Biobank dataset follow established research guidelines. These SNPs, derived from a different random sample of GWAS, have highly significant correlations with BMI. Therefore, provided that the instrument is correctly specified and sufficiently predictive of the endogenous variable, it is expected to demonstrate a strong association with obesity exclusively. Theoretically, Imbens and Rubin (2015) argue that if an instrument is derived from a different random genetic assignment, it is plausible to assume that it remains uncorrelated with the error term in the outcome equation. Given this empirical and theoretical support, the exclusion restriction is likely to hold in this context, minimising concerns about violations of this assumption. Among working-age individuals in our sample, we find a positive and statistically significant association between GRS and obesity (*section C.2 in Appendix C*).

To further assess the plausibility of the exclusion restriction, we regress our IV, the GRS, on the set of control variables used in the main analysis, both individually and jointly. While this procedure does not constitute a formal test of the exclusion restriction, which concerns unobserved confounders rather than observed covariates, it provides a useful check. As shown in Table C.2, the coefficients on the covariates are small in magnitude and, jointly, explain very little of the variation in the instrument ($R^2 < 0.002$). Similarly, the comparison between the full model and the no-control model in Table C.3 demonstrates that the estimated effect of obesity on employment remains stable. The difference between the two models is not statistically significant (t = 1.03), and the corresponding t-statistic confirms that the coefficients in columns 1 and 2 can be considered equivalent. Together, these findings offer additional reassurance that the instrument is not systematically related to observed confounders, supporting the plausibility of the exclusion restriction.

This study focuses on a homogeneous sample of 'White British' participants, consistent with previous research. Similarly to the approach used by Böckerman et al. (2019), the GRS for BMI is applied within a single ethnic group, as different racial or ethnic populations may require distinct combinations of genetic variants to accurately construct a GRS. Furthermore, previous studies have shown that MR remains valid as long as the genetic variants used are not correlated with other characteristics that influence employment outcomes at the population level (Davey Smith et al., 2007; Lawlor et al., 2008a,b; von Hinke et al., 2014).

[7]This condition is also referred to as the independence assumption in some strands of the prior literature (Lee and Lee, 2025).

Relevance of the Instrument

Several studies have successfully included GRS in economic research to examine the causal effects of obesity. For example, Campbell et al. (2021) and Böckerman et al. (2019) applied GRS as an instrument to analyse the impact of obesity on employment, finding significant negative effects on both years of employment and wages. These findings align with the broader literature on the economic consequences of obesity (Böckerman et al., 2019; Cawley, 2004; Morris, 2007; Norton and Han, 2008), further reinforcing the validity of GRS as an instrument.

The instrument must be relevant, meaning it should be correlated with the endogenous treatment variable ($cov(GRS, Obesity)) \neq 0$) and has sufficient variation. The relevance of GRS as an instrument is well-documented in genetic studies. Destri et al. (2022), Yengo et al. (2018), and Locke et al. (2015) established strong associations between GRS for BMI and obesity-related traits, validating its use as an IV. The first-stage results from our analysis indicate a very strong correlation between obesity and the GRS of high BMI.

Strength of the Instrument

Even if the exclusion restriction and relevance assumptions hold, the instrument should not be weak, as weak instruments lead to imprecise and potentially biased causal effect estimates. To assess instrument strength, a Wald F-test is applied. A value potentially exceeding 10 suggests that the instrument is sufficiently strong (Staiger and Stock, 1997), though more recent studies argue for a higher threshold at nearer 100 (Lee et al., 2022). A weak instrument can result in bias toward LPM estimates and lead to incorrect inference due to imprecision. We assess the strength of the instrument using a Wald F-test, with F-statistics of approximately 14,000, well above conventional thresholds for weak instruments (Staiger and Stock, 1997). This provides strong empirical support for valid instrumental variable estimation.

Monotonicity

The monotonicity assumption is a requirement for the identification and interpretation of the Local Average Treatment Effect (LATE) and Local Instrumental Variables (LIV) within the principal stratification framework. This assumption stipulates that the instrumental variable GRS must influence obesity in a consistent direction across all individuals in the population, that is, it cannot increase the likelihood of treatment for some individuals while simultaneously decreasing it for others (Lee and Lee, 2025; Von Hinke et al., 2016).

3.4.3 Marginal Treatment Effects

The Marginal Treatment Effects (MTE) framework allows for the estimation of treatment effects at different points along the distribution of unobserved factors that influence an individual's likelihood of selecting into treatment (Basu, 2014; Basu et al., 2018). In other words, it captures how the impact of obesity varies depending on an individual's likelihood to be living with obesity, which is not directly observable. This is particularly useful in cases where individuals self-select into treatment based on characteristics that are unmeasured but systematically affect their employment outcomes. To illustrate, we adopt a potential outcomes framework as follows:

$$Y_1 = \mu_1(X_O, X_U, \eta) \text{ and } Y_0 = \mu_0(X_O, X_U, \eta), \tag{3.3}$$

Where Y_1 and Y_0 represent the potential employment outcomes for individuals with and without obesity, respectively. X_O is a vector of observed sociodemographics factors as mentioned before, while X_U denotes a vector of unobserved confounders that are believed to also affect employment via treatment selection (obesity). η represents an unobserved random variable that captures all remaining unobserved factors. With the following conditions:

$$(X_O, X_U) \perp\!\!\!\perp \eta \quad \text{and} \quad X_O \perp\!\!\!\perp X_U,$$

where $\perp\!\!\!\perp$ describes statistical independence. The probability of obesity is modelled as:

$$Obesity = 1 \quad \text{if} \quad \mu_{obesity}(X_O, GRS) - U_{Obesity} > 0, \tag{3.4}$$

where $\mu_{Obesity}(\cdot)$ denotes an unknown function with respect to obesity (*Obesity*) and GRS, and $U_{Obesity}$ represents unobserved factors influencing obesity selection, assumed to be correlated with X_U and all other unobserved factors in that also affecting obesity, while $U_{Obesity} \perp\!\!\!\perp \eta$.

Equation (3.3) and Equation (3.4) represent the nonparametric models that conform to the Angrist and Imbens (1995) independence and monotonicity assumptions needed to interpret IV estimates in a model of heterogeneous returns. As in Heckman and Vytlacil (1999), we can rewrite (3.4) as:

$$Obesity = 1 \quad \text{if} \quad P(X_O = x_{Oi}, GRS = grs_i) > V, \tag{3.5}$$

Where $V = F_{U_{Obesity}}[U_{Obesity}|X_O = x_{Oi}, GRS = grs_i], P(x_{Oi}, grs_i) = F_{U_{Obesity|x_{Oi},grs_i}}[\mu_{Obesity}(x_{Oi}, grs_i)]$, and F represents a cumulative distribution function. The distribution of $U_{Obesity}$ conditional on X_O and GRS, by definition $V \sim$ Uniform[0,1].

Following Heckman and Vytlacil (1999), the MTE is identified as:

$$\frac{\partial E(Y \mid X_O = x_{Oi}, GRS = grs_i)}{\partial P(x_{Oi}, grs_i)} = E(Y_1 - Y_0 \mid X_O, V = v_i) = \text{MTE}(x_{Oi}, v_i), \tag{3.6}$$

Where $Y = Obesity \times Y_1 + (1 - Obesity) \times Y_0$ and $v = P(x_{Oi}, grs_i)$. This shows the effect for an individual who is at the margin of choice whether to be with obesity or not such that one's levels of X_O and GRS are just balanced by one's level of V (which includes X_U).

3.4.4 Person-Centered Treatment Effects

The Person-Centered Treatment Effects (PeT) effects can be seen as a personalised, weighted aggregation of MTEs. For any individual, it captures the relevant margins of treatment effect heterogeneity based on their observed characteristics X_O, $P(GRS)$. Hence, the PeT effects (Basu, 2014) can be computed as:

$$\begin{aligned} E_{X_U|X_O,P(GRS),Obesity} E(Y_1 - Y_0 \mid x_{Oi}, P(grs_i), Obesity = 1) &= E((Y_1 - Y_0) \mid x_{Oi}, V < P(grs_i)) \\ &= \frac{\int_0^{P(grs)} \text{MTE}(x_{Oi}, v_i)dv}{P(grs_i)}. \end{aligned} \tag{3.7}$$

Aggregating over the population yields key policy-relevant parameters:

$$ATE = \int_0^1 MTE(x_{Oi}, v_i)dv. \tag{3.8}$$

$$TT = \int_0^{P(grs_i)} MTE(x_{Oi}, v_i)dv. \tag{3.9}$$

$$TUT = \int_{P(grs_i)}^1 MTE(x_{Oi}, v_i)dv. \tag{3.10}$$

For any given individual, it identifies the relevant margins based on their specific values of X_O, $P(GRS)$, and $Obesity$, averaging the MTEs over those margins. Thus, the PeT effects method serves as the observable covariates (X_O and GRS) conditional effect on the treated for those receiving treatment and on the untreated for those not undergoing treatment.

In the absence of essential heterogeneity, the PeT effects simplify to an individualised CATE, where conditioning is applied to the individual's entire observed covariate vector. We focus on explaining the concept of PeT effects and their significance within this analytical framework. This framework provides an understanding of how obesity influences employment, accounting for individual heterogeneity and unobserved selection factors.

Intuition behind using PeT effects

A local instrumental variable approach can be applied to address the challenge of treatment effect heterogeneity when a continuous instrument is available. Unlike binary instruments, continuous instruments, such as the genetic risk score (GRS) used in this study, allow for the identification of marginal treatment effects (MTEs) across different points of the unobserved distribution influencing treatment selection.

In the context of this research, the MTE captures the effect of obesity on employment for individuals whose decision to have obesity or not is finely balanced between observed characteristics (such as GRS, age, sex, deprivation, and education) and unobserved confounders (e.g., individual motivation or lifestyle factors). Conceptually, the MTE is estimated by comparing the outcomes (employment status) of two individuals who are identical in their observed characteristics but differ infinitesimally in their GRS values, say at GRS and GRS+ε. If GRS is a valid strong instrument and exogenous to unobserved determinants of employment, then any difference in employment outcomes must result from a change in treatment status (obesity), induced solely by this small change in GRS.

Because the difference in GRS is infinitesimal, the only individuals affected are those who were marginally indifferent between having obesity or not. Thus, for this margin of individuals, their observed and unobserved characteristics must be in equilibrium. By repeating this comparison at multiple points along the distribution of GRS (e.g., GRS$'$, GRS$' + \varepsilon$, etc.), we can trace out a full schedule of MTEs that vary across the distribution of unobserved confounders, conditional on fixed observed covariates.

Local Instrumental Variable (LIV) methods estimate this schedule by first constructing a control function that models the outcome (employment) as a function of the observed covariates, the instrument-driven propensity to receive treatment, and interactions and nonlinearities of these variables. The partial derivative of the outcome with respect to the IV-based propensity score then captures the MTE at a given point in the unobserved distribution.

Once MTEs are estimated across the entire distribution of unobserved heterogeneity, they can be aggregated to recover standard treatment effect parameters such as the Average Treatment Effect (ATE), the Treatment Effect on the Treated (TT), the Treatment Effect on the Untreated (TUT), and Conditional ATEs. Furthermore, the framework enables the computation of Person-centered Treatment (PeT) effects, which reflect the expected treatment effect for each individual, conditional on both their observed characteristics and their observed treatment status.

The PeT effect for an individual integrates over a distribution of unobserved factors consistent with their actual treatment choice, e.g., an individual with obesity and without an academic degree. For such individuals, the PeT effect is derived by averaging the relevant MTEs over the region of the unobserved distribution consistent with individuals with obesity, conditional on their observed characteristics. This allows us to identify and characterise meaningful treatment effect heterogeneity. For instance, that individuals without an academic degree face a greater employment penalty from obesity than their more-educated peers.

Using the LIV framework and PeT methodology helps us gain a nuanced understanding of how the effect of obesity on employment varies across the population, capturing not only average effects but also the distributional consequences of treatment heterogeneity.

Identification Strategy- PeT Effects

The included continuous IV provides exogenous variation in obesity status, while being uncorrelated with unobserved factors that directly affect employment. This allows us to estimate PeT effects using the following steps:

1. Estimate the first-stage regression by regressing obesity status ($Obesity$) on observed covariates (X_O) and the genetic risk score (GRS) using a probit model. Predict the propensity score, $\hat{P}(x_{Oi}, grs_i)$, for each individual.

2. Ensure that $\hat{P}(x_{Oi}, grs_i)$ has mass at every value (rounded to 0.01) for both individuals with and without obesity. We also verify whether any predicted probabilities are too extreme (i.e., close to 0 or 1), which could indicate poor overlap. In our case, this is not a concern, as the estimated propensity scores range from 0.188 to 0.771.

3. Define $\min_p = \min\{\hat{P}(x_{Oi}, grs_i)\}$ and $\max_p = \max\{\hat{P}(x_{Oi}, grs_i)\}$.

4. Determine the appropriate specification for the second-stage LIV model for probability of being in a paid employment (Y_i) using a probit specification:

$$Y_i = \Phi(\beta_0 + \beta_1 X_{Oi} + \beta_2 \hat{P} + \beta_3 (\hat{P} \cdot X_{Oi}) + \beta_4 (\hat{P} \cdot X_{Oi})^2), \tag{3.11}$$

   Where Φ is a cumulative normal distribution function, and P is $P(x_{Oi}, grs_i)$ in equation (3.11).

5. Specify the estimated LIV, as below:

   (a) Run the second-stage LIV estimation for our outcome, applying a first-degree polynomial in the estimated propensity score, which was selected as the best-fitting model based on goodness-of-fit tests.

   (b) Draw 1,000 replications of $u \sim U(\min_p, \max_p)$.

   (c) Perform numerical integration by computing $\frac{dYi(\cdot)}{d\hat{P}}$, where Y_i is included from equation (3.11), replacing $\hat{P}(x_{Oi}, \mathrm{grs}_i)$ with each value of u_i, and generating 1,000 values for each individual i.

   (d) Compute $Obesity^* = \Phi^{-1}(\hat{P}(x_{Oi}, grs_i))$, also generating 1,000 values for each individual i.

   (e) Compute the PeT by averaging $\frac{dY_i(\cdot)}{d\hat{P}}$ over values of u_i for which $Obesity^* > 0$ if $Obesity = 1$; otherwise, average $\frac{dY_i(\cdot)}{d\hat{P}}$ over values of u_i for which $Obesity^* \leq 0$ if $Obesity = 0$.

6. Specify the control function approach for estimating both the MTE and PeT effects using a probit model. The control function includes a first-degree polynomial in the estimated propensity score. However, unlike linear models, a first-degree polynomial within a nonlinear model does not preclude the presence of essential heterogeneity in the additive scale (Basu, 2011).

7. Compute mean treatment effect parameters by aggregating PeT estimates:

   - Averaging PeT across all observations gives the **ATE**.
   - Averaging PeT for $O = 1$ yields the **TT**.
   - Averaging PeT for $O = 0$ gives the **TUT**.
   - Averaging PeT for X_O (age, degree, Townsend deprivation score, and gender-specific groups), yields the MTEs of X_O.

We explore the underlying mechanisms driving the relationship between obesity and employment for both treated and untreated individuals (with obesity vs. without obesity). In addition, we analyse variations in the estimated effects across different subpopulations, particularly by examining whether the employment consequences of obesity differ by gender, degree, and deprivation.

3.5 Results

We begin by presenting summary statistics for the key variables used in this study, as shown in *Table 3.1*. The sample consists of 284,270 observations. The mean employment rate is 75.5%, indicating a relatively high proportion of individuals in paid employment in our sample of working-age individuals. The prevalence of obesity is 24.1%, suggesting a similar proportion of individuals classified as an individual with obesity in the sample compared to the UK population. The average age in the sample is approximately 53 years, ranging from 38 to 64 years. The Townsend Deprivation score, which measures socioeconomic status, has a mean of -1.45, with a standard deviation of 3, implying that many individuals in the sample live in relatively less deprived areas. 35% of the sample holds an academic degree. This appears to be a feature of the Biobank data, as it is a voluntary sample and tends to attract educated participants.

Table 3.1 Summary Statistics

Variable	**Observations**	**Mean**	**Std. Dev.**	**Min**	**Max**
Employed	284,270	0.755	0.430	0	1
Age	284,270	52.840	6.599	38	64
Townsend Deprivation Index	284,270	-1.454	2.991	-6.26	10.88
Degree	284,270	0.35	0.48	0	1
Female	284,270	0.459	0.498	0	1
Waist (Female)	130,392	83.552	12.67	41.4	171
Waist (Male)	153,878	96.749	11	20	197
BMI	284,270	27.383	4.822	12.121	74.683
Obesity(BMI)	284,270	0.241	0.428	0	1
Obesity(BMIFemale)	130,392	0.223	0.416	0	1
Obesity(BMIMale)	153,878	0.255	0.436	0	1
Obesity (Waist)	284,270	0.554	0.497	0	1
Obesity (WaistToHeight)	284,270	0.199	0.399	0	1
GRS	284,270	-0.206	0.980	-4.709	4.051
Pscore	284,270	0.241	0.111	0.188	0.771

The proportion of females is 46%, close to an even distribution. Lastly, the GRS for high BMI values has a mean of -0.21 with a standard deviation of 0.980, ranging from -4.71 to 4.05, indicating variation in genetic predisposition to obesity across individuals.

Table 3.2 shows that individuals with obesity tend to have lower employment rates (72.3% vs 76.5%) and are less likely to hold a degree (26.3% vs 37.8%) compared to their non-obese counterparts. They are also slightly older on average (53.3 vs. 52.7 years), more socioeconomically deprived (as indicated by higher Townsend scores), and exhibit higher genetic risk scores for obesity, as one would expect. The higher average propensity score in the individuals with obesity group (0.292 vs 0.225) reflects a higher predicted likelihood of obesity, based on observed characteristics and instrumental variables.

Table 3.2 Descriptive Statistics, by Obesity Status(BMI)

	Obesity = 0			Obesity = 1		
Variable	Mean	SD	N	Mean	SD	N
Employed	0.765	0.424	215,792	0.723	0.448	68,478
Age	52.695	6.637	215,792	53.300	6.458	68,478
Townsend	-1.591	2.926	215,792	-1.021	3.149	68,478
Degree	0.378	0.485	215,792	0.263	0.440	68,478
Female	0.469	0.499	215,792	0.425	0.494	68,478
Standard GRS	-0.329	0.958	215,792	0.183	0.944	68,478
Propensity Score	0.225	0.104	215,792	0.292	0.117	68,478

3.5.1 Measures of Obesity

As obesity is derived from the BMI variable, there is a potential for measurement error. BMI is the most commonly used metric for assessing weight status, including classifications of underweight, overweight, and obesity. However, it may sometimes misclassify individuals, leading to inaccuracies in identifying obesity. For instance, individuals with a high proportion of lean muscle and low body fat can be erroneously classified as individuals with obesity, since BMI accounts for total body weight without distinguishing between different components of body composition. To address this, obesity is sometimes measured using waist circumference, which provides an additional indicator of central adiposity. According to clinical guidelines, individuals with waist circumference exceeding 102 cm for males and 88 cm for females are considered at higher risk for obesity-related health complications (NHLBI, 2024).

To evaluate the validity of our obesity measure, we construct two alternative indicators. First, we create a binary (dummy) variable for obesity based on the gender-specific thresholds mentioned in the previous paragraph. Second, we construct an alternative obesity measure based on the waist-to-height ratio, which accounts for the natural variation in waist size due to height, while also controlling for individuals' biological gender. Taking into account height, this ratio offers a more reliable assessment of obesity. Following existing health guidelines mentioned in NICE (2025), we use a threshold of 0.6 or higher to indicate an increased health risk.

We present a summary of the waist circumference data from the UK Biobank in *Table 3.1*, which shows that the average values exceed the NHS defined thresholds for both males and females.[8] The prevalence of obesity based only on waist circumference exceeds 50%, which is more than double the rate when obesity is defined using BMI alone. However, when we incorporate height and define obesity as a waist-to-height ratio greater than 0.6, we find a more realistic prevalence of approximately 20%.

To explore the relationship between these alternative measures and BMI-based obesity, we regress obesity measured by waist-to-height ratio on the BMI-defined obesity variable. The results, presented in *Table 3.3*, indicate a statistically significant and positive association between obesity based on BMI and waist-based measures. Notably, yielding an R-squared of approximately 49%, suggesting a strong correlation.

Table 3.4 reveals notable patterns in the classification of obesity based on waist circumference and BMI. 206,753 individuals are classified as without obesity and 47,497 classified as individuals with obesity, by both measures. However, there are discrepancies: 20,981 individuals are categorised as with obesity based on waist circumference ratio, but not by BMI. These discrepancies may arise due to differences in body composition, muscle mass, or limitations in waist circumference as a measure of adiposity. In addition, if an individual is classified as with obesity based on BMI, their waist circumference generally aligns with this classification, with only a small number of exceptions (9,039 out of 284,270 individuals). A total of 30,020 individuals out of the total sample of 284,270 have different obesity classifications depending on the measurement method used, while the remaining individuals are consistently classified across both methods (approximately 89.4%).

We present a cross-tabulation in *Table 3.4* to assess the degree of agreement between the BMI-based and waist-to-height ratio obesity classifications. A Chi-squared test of independence strongly rejects the null hypothesis of no association between the measures ($\chi^2(1)$ = 140,000, $p < 0.001$), indicating a statistically significant relationship between the two. To further evaluate the agreement beyond chance, we calculated Cohen's Kappa,

[8] According to the most recent estimates, the UK obesity rate is 26.2% (NHS, 2023), compared to the 2008 estimates of 24% for females and 25% for males (NHS Digital, 2009).

Table 3.3 Regression of Waist-to-Height Ratio-Based Obesity on BMI-Based Obesity

	Obesity(BMI)
Obesity (WaistToHeight)	0.748^{***}
	(0.0014)
Constant	0.092^{***}
	(0.0006)
Observations	284,270
R-squared	0.488
Adj. R-squared	0.488
Root MSE	0.306
F-statistic	>99,999
Prob > F	0.0000

*** Significant at the 1% level

which yielded a value of 0.693. This represents substantial agreement between the measures according to the standard interpretation of Kappa values (of 0.69).

Table 3.4 Cross-tabulation of obesity by waist circumference and BMI-based obesity measure

	Obesity(WaistRatio)		
Obesity (BMI)	0	1	Total
0	206,753	9,039	215,792
1	20,981	47,497	68,478
Total	227,734	56,536	284,270

Pearson Chi-squared (1):	140,000	($p < 0.001$)
Observed Agreement:	89.44%	
Expected Agreement:	65.61%	
Cohen's Kappa:	0.693	($p < 0.001$)

Nevertheless, we do not adopt waist circumference as our primary measure of obesity. Our analysis is based on an IV approach that uses BMI as the measure of obesity, with the genetic risk score as the instrument. As a robustness check, we conduct an additional analysis using alternative measures of obesity, applying the same estimation methods to assess whether the results remain consistent (*Appendix C*).

3.5.2 Linear Approach

The results in *Table 3.5* compare the estimated effects of obesity on paid employment across different linear regression models, LPM and Mendelian randomisation (MR). Each method provides a distinct perspective on the relationship between obesity and employment status, while MR accounts for the potential endogeneity of obesity.

Table 3.5 Comparison of Obesity Effect on Employment Across LPM and MR (2SLS)

Variable	**LPM**	**MR (2SLS)**
Obesity	-0.018***	-0.042***
	(0.0018)	(0.0082)
Age	-0.021***	-0.021***
	(0.0001)	(0.0001)
Townsend	-0.011***	-0.011***
	(0.0003)	(0.0003)
Degree	0.042***	0.040***
	(0.0016)	(0.0017)
Female	-0.001	-0.002
	(0.0016)	(0.0016)
Constant	1.853***	1.858***
	(0.0066)	(0.0068)
Observations	284,270	284,270
R-squared	0.1145	0.1139
Weak identification test (Cragg-Donald F: First-stage)	-	14,000
Endogenous variable	Obesity	
Instrumental variable	Genetic risk score for BMI	

Note: Dependent variable is being in paid employment. *** Significant at the 1% level.

The LPM results (column 1) suggest a negative association between obesity and paid employment, with a coefficient of -0.018 that is statistically significant at the 1% level. This implies that, ceteris paribus, obesity is associated with a 1.8 percentage point reduction in the likelihood of being in paid employment. However, LPM estimates may be biased due to unobserved confounders that correlate with both obesity and employment. These effects may bias the causal effect of obesity on employment status.

The MR results presented in *Table 3.5*, which leverage a genetic instrumental variable to address endogeneity, yield a much greater negative effect of obesity on paid employment (more than twice greater, at a reduction of 4.2% point), statistically significant at the 1%

level. This increase in magnitude suggests that the LPM estimates may be understated due to omitted variable bias. By isolating the variation in obesity attributable to genetic factors, the MR approach provides a more robust estimation of the causal impact of obesity on employment. Other sociodemographic factors associated with the likelihood of being in paid employment exhibit a pattern consistent with the findings of the LPM analysis.

To assess the relevance assumption in the IV approach, as discussed in *subsection 3.4.2*, we examine the first-stage F-statistic to determine whether the IV is weak or strong. As shown in *Table 3.5*, the first-stage F-statistic is 14,000, which is substantially above the commonly accepted threshold of 10 (Staiger and Stock, 1997), as well as the more conservative benchmark of 100 suggested in recent research (Lee et al., 2022). This provides strong evidence that the IV is relevant and satisfies the relevance assumption.

3.5.3 Marginal Treatment Effect

The results presented in *Table 3.6* show the MTE parameters estimated by the local instrumental variable approach and indicate significant effects for several variables in the treated and untreated groups. The coefficient on obesity indicates a decrease of 4 percentage points in the probability of the outcome, with a standard error of 0.008, suggesting a statistically significant and substantial negative association. Variables such as an academic degree and gender, also exhibit significant effects with some differences between treated and untreated groups. These effects, as well as the impact of IV, can be further visualised in *Figure 3.3*, which provides a graphical representation of a nonlinear relationship between obesity and the genetic risk score for obesity (higher BMI), where the dotted line represents the ATE (reduction of 4 percentage points), suggesting the presence of nonlinearity and a heterogeneous effect of GRS on the likelihood of having obesity.

The curve in *Figure 3.3* is downward-sloping in the lower quantiles of unobserved factors influencing obesity selection, indicating that individuals with a lower ratio of unobserved characteristics affecting obesity decision, tend to experience less negative treatment effects on employment. In contrast, individuals with a higher degree of unobserved factors affecting obesity face more negative employment penalties if they do have obesity.

3.5.4 Person-centered Treatment Effect

Figure 3.4 illustrates the distribution of propensity scores for both treated and untreated observations. A total of 38 individuals fall outside the region of common support and are excluded from the analysis. This ensures that comparisons between groups are made within

Table 3.6 Marginal Treatment Effect Results

Variable	Treated	Untreated
Obesity ATE	-0.039***	-
	(0.008)	-
Age	-0.022***	-0.021***
	(0.0002)	(0.0002)
Townsend	-0.018***	-0.008***
	(0.001)	(0.0003)
Female	-0.011***	0.003*
	(0.003)	(0.002)
Degree	0.060***	0.036***
	(0.003)	(0.002)
k	-0.016***	-0.013**
	(0.006)	(0.005)
Constant	1.846***	1.851***
	(0.014)	(0.009)
Mills: $\rho_1 - \rho_0$	-0.003	-
	(0.007)	-
Number of observations: 284,270		
Treatment model: Probit		

a region where treatment and control observations overlap, thus improving the credibility of the estimated PeT effects and satisfying a key condition for causal inference.

The results presented in *Table 3.7* highlight a negative relationship between obesity and paid employment, as indicated by the estimated PeT effects. The negative mean values observed across the entire sample, as well as the treatment effects for both the treated and non-treated groups, suggest that individuals with higher levels of obesity are less likely to be in any form of paid employment. The results of the PeT effects indicate that the impact of obesity on the probability of having paid employment is more closely aligned with the MR estimates than with the LPM estimates used in this study. However, the results show that the effect is slightly marginally greater compared to MR analysis (0.043 vs 0.042). This discrepancy may arise from the exclusion of a potential heterogeneity of unobserved factors in the MR analysis, which could influence the probability of living with obesity.

Table 3.8 presents the distribution of the person centred treatment effect for individuals with and without an academic degree. The ATE for individuals with a degree is 0.015, which is different compared to those without a degree, whose effect is -0.051. This suggests that individuals with an academic degree have a positive effect of obesity on the probability of

Fig. 3.3 Nonlinear Effects of Obesity and GRS

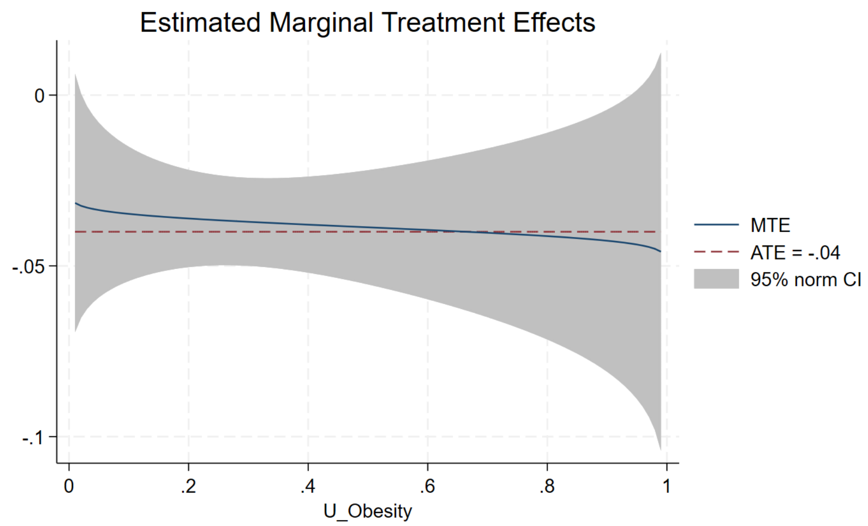

Note: $U_{Obesity}$ represents unobserved factors influencing obesity selection.

Fig. 3.4 Propensity Score, By Obesity Status

being in a paid employment. This indicates a significant difference in the effect of having an academic degree on employment status.

Table 3.8 also shows that ATE differs by area-level deprivation. Specifically, individuals living in more deprived areas exhibit a larger negative association between obesity and employment, while those in less deprived areas show a positive ATE. The TT and TUT estimates follow a similar pattern.

Table 3.7 Effect of Obesity on Paid Employment

	Full Sample	**By Obesity Status**	
	ATE	TT	TUT
Obesity	-0.028	-0.043	-0.023
SD	0.135	0.141	0.133
Observations	284,270	68,478	215,792

Notes: The table reports the mean of *pet_employed* across the full sample, by obesity status.

Table 3.8 Comparison of PeT Effects by Degree and Deprivation Status

By Degree Status		
	Degree = 1	**Degree = 0**
ATE	0.015	-0.052
SD	0.112	0.142
TT	0.016	-0.069
SD	0.115	0.146
TUT	0.015	-0.045
SD	0.112	0.140
Observations	99,570	184,700
By Deprivation Status		
	Deprived = 1	**Deprived = 0**
ATE	-0.203	0.036
SD	0.118	0.069
TT	-0.209	0.031
SD	0.117	0.072
TUT	-0.201	0.038
SD	0.118	0.068
Observations	76,517	207,753

Note: PeT effects are reported as ATE, TT, and TUT for degree and deprivation. Deprivation is based on the Townsend Deprivation Index (values > 0 = deprived).

3.5.5 Heterogeneity by Gender

As the primary analysis demonstrates, the LPM model is inaccurate and biased due to endogeneity issues. To address this, we apply the MR model to gender-specific subsamples, incorporating the PeT effects to examine gender-based differences. The results in *Table 3.9* consistently show a negative relationship between obesity and paid employment across both genders in the MR approach. The negative coefficients for obesity indicate that higher levels of obesity are significantly associated with a lower likelihood of being in paid employment. The effect is significantly greater for females compared to males, indicating a stronger weight-related penalty for women, as expected. This effect is statistically significant at the 1% level.

When the PeT effects approach is applied, the impact of obesity on the probability of being in paid employment becomes significantly more pronounced for males. Among females, the estimated effect of obesity on employment decreased from a 4.7 percentage point reduction to 2.1 percentage points. In contrast, for males, the negative effect increases significantly in magnitude (-0.066 vs. -0.039), indicating a stronger association between obesity and reduced employment prospects. In addition, the point estimates for the Average Treatment Effect (ATE) and Treatment on the Untreated (TUT) effects also diverge, suggesting differential impacts of obesity on employment outcomes between treated and untreated individuals.

Table 3.9 Effect of Obesity on Paid Employment by Gender

	Females	**Males**
Panel A: MR Estimates		
Obesity	-0.047***	-0.039***
	(0.012)	(0.011)
Panel B: PeT Effects		
ATE	-0.001	-0.051
	0.123	0.142
TT	-0.021	-0.066
	0.134	0.147
TUT	0.004	-0.046
	0.119	0.141
Observations	130,387	153,871

Notes: Estimates in Panel A are derived using Mendelian Randomisation (MR). PeT effects in Panel B correspond to the Average Treatment Effect (ATE), effect on the Treated (TT), and effect on the Untreated (TUT). Standard deviations are reported in Panel B. *** $p < 0.01$.

Figure 3.5 displays PeT effects on paid employment by age and gender, revealing a notable age-dependent trend: PeT effects increase with age for both genders, with a steeper rise among females. Although the figure suggests that older females experience stronger positive effects than males, the overall mean PeT effect is closer to zero for females than for males (*Table 3.9*), likely reflecting a concentration of females in younger age brackets with smaller or more negative effects.

Fig. 3.5 Dependence of PeT Effect on Age by Gender

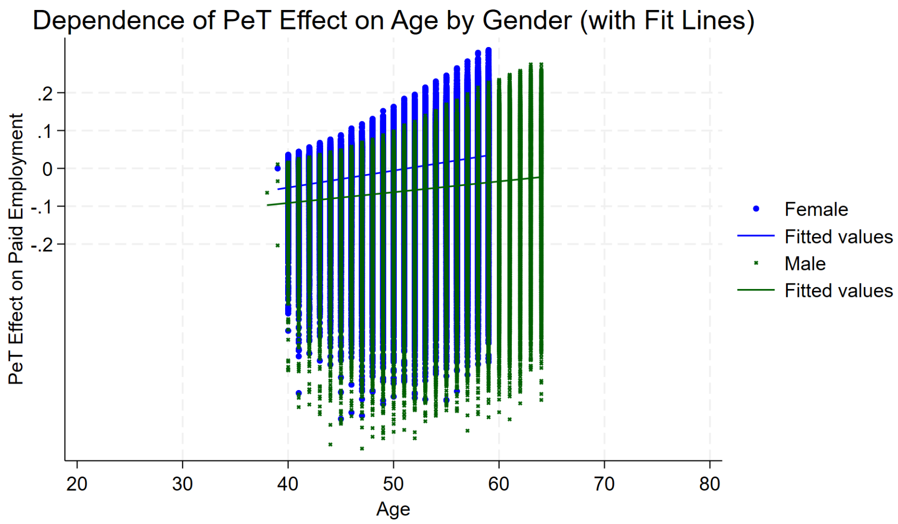

Since the MR average effects of the observed covariates, particularly gender, differ from the ATEs estimated by the LIV approach, we find compelling evidence of essential heterogeneity in the likelihood of being in paid employment.

3.6 Discussion and Conclusion

The relationship between obesity and paid employment is complex, and this study contributes to a deeper understanding of this relationship by applying a variety of econometric techniques to address potential sources of bias. The significant discrepancy between Linear Probability Model (LPM) and Mendelian Randomisation (MR) results suggests that the negative association observed in the LPM approach may be influenced by confounding factors, rather than reflecting a direct causal relationship between obesity and employment outcomes. For example, unobserved variables, such as general health status or discrimination within the labour market, can simultaneously affect both obesity and employment outcomes, thus inflating the LPM coefficient. This finding is consistent with previous literature that underscores the potential role of unobserved factors, such as motivation, in shaping both obesity and employment outcomes (Böckerman et al., 2019; Cawley, 2004; Kesaite and Greve, 2024; Morris, 2007).

The use of MR analysis, a more rigorous approach to overcome endogeneity, reveals that the true causal effect of obesity on paid employment is likely greater than the LPM estimates suggest. By using genetic variants as instrumental variables, MR mitigates the risk of reverse causality, where the relationship between obesity and employment could be bidirectional. This is particularly important in light of previous studies that have reported mixed findings, with some showing that obesity leads to lower employment rates (Böckerman et al., 2019; Caliendo and Gehrsitz, 2016; Cawley, 2004; Kesaite and Greve, 2024), while others suggest that employment status may influence the likelihood of obesity (Caliendo and Gehrsitz, 2016; Morris, 2007).

The subsample analysis reveals heterogeneous effects by gender, suggesting that gender discrimination remains a factor in employment outcomes, alongside differences in obesity prevalence (25.5% for males and 22.3% for females in our sample). The effect of obesity on employment differs by gender, with a more pronounced negative impact observed for women, when using linear estimation methods. This is consistent with numerous studies highlighting persistent gender disparities in employment and earnings, as well as a greater weight-related penalty for females (Cawley, 2004; Goldin, 2014; Tyrrell et al., 2016).

Applying the nonlinear models (MTE and PeT effects) provides deeper insights into the complex negative relationship between obesity and employment, particularly when accounting for the effects of other unobservered factors that can affect the decision of having obesity, including factors such as job satisfaction, work motivation, and ability, which may influence both the likelihood of having obesity and employment outcomes. These factors are likely to vary between individuals, highlighting the necessity of further estimations to evaluate the robustness of alternative causal models. Notably, the MTE and PeT effect estimates follow around those derived from MR, with reporting marginally different estimates. By addressing these unobserved factors, the PeT effects framework offers a more refined understanding of how obesity influences labour market participation, as failing to do so may lead to inaccurate conclusions about both the determinants of obesity and its impact on employment probabilities (Basu, 2014; Basu et al., 2018).

The findings of the PeT effects approach suggest that the impact of obesity on the likelihood of being in paid employment is greater than that estimated using LPM and MR methods. Specifically, obesity is associated with a 4.3 percentage point reduction in the probability of being in paid employment, compared to reductions of 1.8 and 4.2 percentage points estimated by the LPM and MR approaches, respectively. The consistent negative impact of obesity on employment across various estimation methods, coupled with differences in marginal effect sizes, demonstrates the presence of essential heterogeneity and emphasises the importance of incorporating PeT effects alongside traditional linear models.

The PeT effects approach further reveals a striking contrast: for females, the impact of obesity on employment is halved when accounting for essential heterogeneity, whereas for males, it nearly doubles. This pattern suggests that the labour market penalty associated with obesity may be more pronounced for men. One possible explanation, consistent with the literature, is that gendered norms surrounding physical appearance and perceived productivity differ between men and women, potentially leading to asymmetric labour market penalties associated with obesity. Existing evidence points to mixed gender patterns: Cawley (2004) finds a wage and employment penalty associated with obesity for women but not for men, whereas Dackehag et al. (2015) reports stronger labour market penalties for men living with obesity. These contrasting findings suggest that the role of gender norms and occupational expectations may vary across institutional and labour market contexts. We therefore interpret our results as suggestive and acknowledge that further research is needed to directly identify the mechanisms driving these gender differences.

In summary, the existing literature highlights the non-negligible impact of obesity on employment. Numerous studies have documented the negative association between obesity and factors such as wages, employment probabilities, and absenteeism. However, some of these studies do not address endogeneity, which is critical for establishing causality. This study provides causal estimates of the impact of obesity on employment status using various econometric techniques and a genetic instrumental variable. The consistently negative and statistically significant findings across multiple econometric models underscore the robustness of the results and highlight important policy implications. These findings accentuate the need for targeted interventions to mitigate their adverse effects on employment and economic participation.

Targeted policies that acknowledge gender-specific dynamics are warranted. For example, in the UK, the government's Work, Health and Disability Green Paper (Department for Work and Pensions and Department of Health, 2016) already emphasises the need for tailored support to help people with health conditions enter or remain in employment. Our findings reinforce such initiatives and also highlight the need for gender-responsive approaches to policy design that address the intersection of health and labour market status.

These findings are also relevant beyond the UK, particularly for countries implementing comparable strategies to address obesity within the context of labour market participation. For example, the United States emphasise workplace wellness through its National Strategy on Obesity Centers for Disease Control and Prevention (2024). Similarly, Australia's National

Obesity Strategy 2022–2032 explicitly recognises the role of employment and recommends gender-sensitive approaches Australian Government Department of Health and Aged Care (2022). Our results support these international efforts. Future research should further investigate these mechanisms by applying diverse datasets, including non-volunteer samples.

Discussion

This thesis addresses three key issues at the intersection of health, income, and employment, using empirical data to provide insights into real-world problems. The findings from each chapter have significant policy implications, particularly in areas such as mental health, income inequality, and obesity, which have become critical topics in both public health and economic policy. In this section, we discuss the policy implications of these findings, drawing on existing literature and offering suggestions for cost-effective interventions.

The first chapter reveals a significant relationship between mental health and early retirement decisions, particularly among individuals aged 50-64 in Australia. The study's findings underscore the need for targeted interventions to support individuals in this age group who may be at risk of early retirement due to mental health challenges.

Mental health issues, such as depression and anxiety, are known to contribute to reduced work capacity and early exit from the labour market (Bryan et al., 2022). Research has shown that mental health can act as a barrier to continued employment, with individuals experiencing psychological distress more likely to retire early (Bubonya et al., 2019; Butterworth et al., 2006). Given this, policy efforts could focus on improving mental health resources in workplaces. For instance, enhancing Employee Assistance Programs (EAPs) and offering mental health screenings as part of routine occupational health services would help identify and address mental health issues before they lead to early retirement. Additionally, given the gender differences observed in our findings, it may be beneficial for employers to address mental health stigma with tailored approaches for each gender. Workplace mental health initiatives are not only beneficial for employees' well-being but can also reduce the financial burden on social welfare systems by keeping individuals in the workforce longer (Butterworth et al., 2006).

A cost-effective strategy could be implementing workplace mental health programs that offer counseling, stress management, and resilience training. These programs, when coupled with strong social safety nets and employer-supported mental health initiatives, can alleviate the negative effects of mental health on early retirement decisions.

The second chapter of this thesis examines the effects of both absolute and relative income on health outcomes in Israel, revealing that while higher absolute income improves health, the effects of relative income are more pronounced, especially among religious groups. These findings suggest that addressing income inequality could have significant public health benefits.

Research has long shown that income inequality is a strong predictor of health disparities within populations (Wilkinson, 1996). The "relative deprivation hypothesis" posits that individuals' health outcomes are influenced not only by their own income but also by their income in relation to others (Jones and Wildman, 2008). In Israel, the effects of relative income were found to be particularly significant across different religious groups, suggesting that policies aimed at reducing income inequality could have a positive impact on public health.

A cost-effective policy response could involve the implementation of progressive taxation systems, such as wealth taxes or higher-income tax rates, along with redistribution programs targeted at lower-income groups. Additionally, policies focused on enhancing access to health services for lower-income populations could help mitigate the health disparities observed across religious groups. For example, expanding public health insurance coverage for marginalised groups and investing in community health centers could improve access to care and reduce income-related health disparities.

The third chapter of this thesis highlights the negative relationship between obesity and employment outcomes, showing that obesity reduces the likelihood of being employed in the UK. This finding is consistent with previous studies, which have demonstrated that obesity is associated with lower employment rates, higher absenteeism, and lower productivity in the workplace (Böckerman et al., 2019; Cawley, 2004).

Addressing obesity not only has public health benefits but could also lead to broader economic gains. Policymakers should consider promoting workplace wellness programs that focus on obesity prevention, such as subsidised gym memberships, nutritional counseling, and regular health screenings. Research has shown that workplace wellness programs can lead to improved employee health and lower healthcare costs for employers (Baicker et al., 2010). In addition to workplace programs, policymakers could provide tax incentives to businesses that encourage healthy lifestyles for their employees, such as offering wellness-related benefits.

On a national level, implementing public health campaigns to raise awareness about obesity and its consequences could complement workplace initiatives. Public health programs, such as those that promote healthy eating, physical activity, and mental health, could reduce the societal costs of obesity. In the UK, for example, the government has already implemented programs like the "Sugar Tax" to reduce sugary drink consumption, which can help address the root causes of obesity. Expanding these types of programs could yield significant economic and health benefits in the long run.

Conclusion

This thesis provides a comprehensive empirical investigation into the interconnections between health, income, and employment, offering new insights into how individual health status and socioeconomic conditions shape labour market outcomes. Across the three empirical chapters, evidence emerges consistently: health is not only a consequence of economic position but a critical determinant of economic behaviour and labour force participation. This research contributes to the growing body of work examining the consequences of health disparities and their underlying economic determinants.

The urgency of addressing mental health, socioeconomic inequality, and obesity is intensifying within the broader context of demographic and economic transformation. Mental health disorders, in particular, have become increasingly prevalent in high-income countries. In England alone, approximately 3.8 million individuals are currently in contact with NHS mental health services, a marked increase compared to pre-pandemic levels (NHS England, 2024). Availability to treatment in a timely manner is not only vital for individual wellbeing but also for economic sustainability. Recent estimates suggest that the total annual costs of mental health are around £118 billion in the United Kingdom, largely through the loss of days in work retention, presenteeism, lower productivity, and the costs associated with support from informal carers. (McDaid and Park, 2022). These concerns are particularly salient given the declining stigma surrounding mental illness and the growing precision of diagnostic tools, which together enhance both the feasibility and the necessity of early intervention (Catania et al., 2011).

Similar concerns are evident in Australia, where mental illness accounts for a substantial share of the national burden of disease. According to the Australian Bureau of Statistics, mental and substance use disorders accounted for approximately 13% of Australia's total disease burden in 2022. Furthermore, nearly 43% of the population had experienced a mental illness at some point in their life, making mental health conditions one of the leading contributors to disability-adjusted life years (DALYs) lost in the country (Australian Bureau of Statistics, 2022a). The economic cost is similarly significant. A Productivity Commission report estimated that mental ill-health costs the Australian economy up to AU$70 billion annually in direct economic costs, with further AU$150 billion attributed to poorer health outcomes and reduced life expectancy among individuals living with mental illness (Productivity Commission, 2020). In this light, proactive investment in mental health services represents not only a public health imperative but also an economically prudent strategy.

The proportion of people aged 85 and over is expected to nearly double by 2042 in Australia, placing increasing pressure on pension systems and healthcare services (Australian Bureau of Statistics, 2018). In response, countries like Denmark have already raised the retirement age to 70 for those born after 1970, aligning pension policies with longer life expectancies. However, the average healthy working life expectancy at age 50 is only 9.42 years, suggesting that many may not remain healthy at work until state pension age (Parker et al., 2020).

This growing mismatch between rising state pension ages and limited healthy working life expectancy highlights a critical policy challenge. As the population ages, maintaining both physical and mental health into later life becomes essential, not only to extend individuals' participation in the workforce but also to prevent early exits due to illness or disability. Without targeted health and labour market interventions, many may face an extended period of economic vulnerability before reaching the state pension age, exacerbating income inequality and placing further strain on healthcare and welfare systems. Studies have shown that older workers in poor health are significantly more likely to retire early or transition into disability pensions, particularly in lower-income or less flexible occupations (Butterworth et al., 2006; Disney et al., 2006).

To gain a deeper understanding of these issues, the first chapter of this thesis examines the causal impact of mental health on early retirement decisions. Using rich longitudinal data from Australia and applying a range of econometric techniques, culminating in a discrete-time survival model combined with a two-stage residual inclusion approach. This chapter helps to identify and validate the causal relationship. The findings show that mental health issues significantly increase the likelihood of premature labour market exit, particularly among men. The use of the death of a close friend as an instrument provides useful variation to address endogeneity, and robustness checks confirm the reliability of the mental health measure through comparisons with more objective indicators. These findings lend empirical support to initiatives like the OECD's call for integrating mental health into workforce strategies and reinforce the argument that improving mental health among older workers is not only a public health priority, but also a labour market necessity.

The second chapter explores the dual role of absolute and relative income in shaping health outcomes. To capture the complexity of this relationship, a range of econometric techniques are applied, including a Probit Mundlak approach (Mundlak, 1978), Hausman-Taylor (Hausman and Taylor, 1981) estimator, and a semiparametric approach (Robinson, 1988) that allows for nonlinearity in the income-health gradient, as one would expect. Using multiple estimators enables a more robust understanding of the relationship, as relying on a single model may obscure important nuances. In this context, the use of different estimation strategies also serves as a form of validation, ensuring that the observed effects are not artefacts of model specification. The consistently statistically significant effect of relative income across all specifications reinforces the importance of psychosocial pathways in health inequality, echoing findings in the literature on social comparison and status anxiety (Wilkinson, 1996).

Variation across religious groups further suggests that cultural and social contexts shape how individuals interpret and respond to economic comparisons, thus influencing health outcomes (Abu Riha, 2015; Cooperman et al., 2016). This highlights the need for policies that are not only economically informed, but also locally and culturally sensitive. The findings suggest that the influence of income, whether absolute or relative, is likely mediated by factors such as identity, social norms, and group-specific expectations, which are largely captured through religious affiliation in the uniquely diverse population of Israel. These insights are particularly relevant in an era of widening income inequality and growing recognition of the broader social determinants of health, a focus recently reaffirmed by the WHO's Commission on Social Determinants of Health in 2024 (WHO, 2024).

Another health issue that is particularly important in the current climate is obesity, which continues to pose a significant public health challenge with considerable economic consequences. Innovative interventions, such as providing weight-loss treatments such as Ozempic to unemployed individuals living with obesity in the UK, are being explored to improve employment prospects and reduce healthcare costs. During the COVID-19 pandemic, individuals classified as obese faced a higher risk of infection and severe outcomes, underscoring the urgent need to address this health condition (Diabetes UK, 2024; Popkin et al., 2020).

The increasing prevalence of obesity continues to raise serious public health and economic concerns, particularly with regard to its implications for labour market outcomes. A substantial body of evidence has established a consistent negative association between obesity and employment-related indicators (Baum and Ford, 2004; Busetta et al., 2020; Cawley, 2004; Li et al., 2022; Morris, 2007). These effects are driven by a variety of interrelated mechanisms, including health-related productivity losses, weight-based discrimination in hiring and promotion, and premature exit from the labour force. Taken together, these findings highlight the urgency of integrating obesity prevention and management into broader economic and social policy frameworks, not only to improve health outcomes but also to enhance workforce participation and reduce economic disparities.

Building on this, the third chapter uses novel methods and rich UK Biobank data to identify the causal impact of obesity on labour market participation. Leveraging genetic instruments via Mendelian Randomisation and estimating marginal and person-centred treatment effects, this chapter adds depth to our understanding of heterogeneity in the likelihood of being obese and the returns to employment. The robust negative causal impact of obesity on employment strengthens the case for viewing obesity not only as a health concern, but as a barrier to economic participation. These findings align with the growing international recognition of the productivity costs of obesity, for instance, in recent policies aimed at workplace wellness in the UK and the US, as well as WHO recommendations on addressing obesity through integrated public health and employment strategies. In addition, these findings align with the OECD's recognition of obesity's economic burden and the need for comprehensive prevention policies (Cecchini and Vuik, 2019).

In all three chapters, this thesis demonstrates the analytical value of combining causal inference methods with detailed data to uncover the underlying mechanisms, thus addressing gaps in previous research. The use of instrumental variable strategies (both traditional and genetic), survival models, and treatment effect heterogeneity analysis positions the research at the methodological frontier of applied health and labour economics. It contributes not only to better empirical identification but also to a more nuanced understanding of how health conditions interact with structural inequalities to shape economic outcomes.

The findings of this thesis highlight the growing exigency of addressing mental health, inequality, and obesity within the broader context of demographic and economic change. Obesity and mental health issues are increasingly prevalent with an increasing prevalence around the world. with early obesity and mental health challenges having long-term implications. For example, children with serious mental health issues are significantly more likely to experience limited work capacity in adulthood (IPPR, 2025). As understanding the long-term effect of these issues improved while countries trying to address this by policies, early intervention becomes both more feasible and necessary.

The relative income hypothesis further illuminates the complexities of economic inequality. Individuals' perceptions of their well-being are influenced not only by absolute income levels but also by their income relative to others (Cui and Chang, 2021; Jones and Wildman, 2008). This suggests that policies aimed at reducing inequality must consider social and cultural contexts to be effective. Hence, the interconnected issues of mental health, demographic change, economic inequality, and obesity require comprehensive and forward-thinking policies. Addressing these challenges holistically will be essential for promoting individual well-being and ensuring sustainable economic growth in the coming years.

The consistent finding of gender differences in two chapters of this thesis also warrants further investigation. Why men seem more vulnerable to mental health and obesity-related exits from the labour force and why obesity affects women differently in employment terms are questions with clear implications for targeted policy design.

Finally, this thesis contributes to an emerging consensus that the boundaries between health policy, labour market policy, and social protection are increasingly porous. Effective interventions must cut across these domains. This research supports a call for integrated evidence-based policies that recognise health as a driver, not just an outcome, of economic opportunity.

There are several promising extensions that emerge from the findings and limitations of this thesis, which could help shape future research. First, while the use of self-reported health measures is widespread, it introduces potential concerns about reporting bias. Future studies could strengthen causal inference by incorporating objective measures of health, such as clinical assessments, administrative health records, or biomarker data, alongside survey responses. Second, Chapters 1 and 2 rely on relatively specific information, such as subsamples (e.g., religious groups) or IV analysis. Expanding the number of waves with an increased sample size or applying other validation tests to improve measurement precision would enhance statistical power and allow for more robust analysis. Third, although this thesis draws on data from three different national contexts (Australia, Israel, and the UK), the analyses are conducted independently. A valuable direction for future work would be to implement a comparative cross-country framework that systematically examines how institutional, cultural, and policy environments moderate the health-income-employment relationship.

In Chapter 3, mendelian Randomisation is included. Although powerful, this approach is subject to limitations due to the voluntary nature of participation in biobank studies, which tend to attract individuals from specific socio-economic or educational backgrounds. Future research could explore ways to improve generalisability, although the development of randomised population-based genetic datasets remains a considerable challenge due to ethical and logistic constraints. Capturing these mechanisms through time-use diaries, qualitative interviews, or richer behavioural modules could provide deeper insight into the causal pathways at play. Finally, the growing availability of high-dimensional genetic and longitudinal data presents opportunities to apply advanced econometric, structural, and machine learning techniques. Furthermore, as data availability improves, exploring the interactions between early life disadvantage, genetic predispositions, and adult health and economic outcomes could yield valuable insights. These methods could enhance estimation precision, capture treatment heterogeneity, and generate more accurate out-of-sample predictions, thereby further advancing the evidence base at the intersection of health and labour economics.

References

Aalen, O. O. (1989). A linear regression model for the analysis of life times. *Statistics in Medicine*, 8(8):907–925.

Abu Riha, H. (2015). Religious coping with life crisis: Literature review. *Social Welfare*, 3(3):323–344.

Adams, P., Hurd, M., McFadden, D., Merril, A., and Ribeiro, T. (2003). Healthy, wealthy and wise? tests for the direct causal paths between health and socioeconomic status. *Journal of Econometrics*, 112(1):3–56.

Alderwick, H. and Dixon, J. (2019). The nhs long term plan. *BMJ (Clinical research ed.)*, 364:l84.

Alexandre, P. K. and French, M. T. (2001). Labor supply of poor residents in metropolitan miami, florida: The role of depression and the co-morbid effects of substance use. *Journal of Mental Health Policy and Economics*, 4(4):161–174.

American Psychiatric Association (2013). *Diagnostic and Statistical Manual of Mental Disorders: DSM-5*. American Psychiatric Association, Washington, D.C.; London, 5th edition.

Andersen, S. H., Richmond-Rakerd, L. S., Moffitt, T. E., and Caspi, A. (2024). The causal effect of mental health on labor market outcomes: The case of stress-related mental disorders following a human-made disaster. *Proceedings of the National Academy of Sciences*, 121(27):e2316423121.

Angrist, J. and Imbens, G. (1995). Identification and estimation of local average treatment effects. *Econometrica*, 108(2):245–254.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Apouey, B. H., Guven, C., and Senik, C. (2019). Retirement and unexpected health shocks. *Economics & Human Biology*, 33:116–123.

Australian Bureau of Statistics (2008). National Survey of Mental Health and Wellbeing: Summary of Results, 2007. Canberra: ABS, Cat. no. (4326.0).

Australian Bureau of Statistics (2018). Population projections, australia, 2017 (base) – 2066. https://www.abs.gov.au/statistics/people/population/population-projections-australia/2017-base-2066. Accessed: 03 May 2025.

Australian Bureau of Statistics (2022a). National study of mental health and wellbeing. https://www.abs.gov.au/statistics/health/mental-health/national-study-mental-health-and-wellbeing/latest-release. Published 2020–2022. Accessed: 23 May 2025.

Australian Bureau of Statistics (2022b). National study of mental health and wellbeing-external site opens in new window. https://www.abs.gov.au/statistics/health/mental-health/national-study-mental-health-and-wellbeing/external-site-opens-new-window. Accessed on July 13, 2023.

Australian Government Department of Health and Aged Care (2022). National obesity strategy 2022–2032. https://www.health.gov.au/resources/publications/national-obesity-strategy-2022-2032. Accessed: 2025-04-29.

Australian Institute of Health and Welfare (2021). Mental health services in Australia. Retrieved from https://www.aihw.gov.au/reports/mental-health-services/mental-health-services-in-australia.

Baicker, K., Cutler, D., and Song, Z. (2010). Workplace wellness programs can generate savings. *Health Affairs*, 29(2):304–311.

Bajorek, Z. and Bevan, S. (2019). Obesity and work: Challenging stigma and discrimination. Retrieved 1 May 2025.

Baker, C. and Kirk-Wade, E. (2023). Mental health statistics: prevalence, services and funding in england. *House of Commons Library*.

Balan, T. A. and Putter, H. (2019). Nonproportional hazards and unobserved heterogeneity in clustered survival data: When can we tell the difference? *Statistics in Medicine*, 38(18):3405–3420.

Bartel, A. and Taubman, P. (1979). Health and labor market success: The role of various diseases. *The Review of Economics and Statistics*, 61(1):1–8.

Bartel, A. and Taubman, P. (1986). Some economic and demographic consequences of mental illness. *Journal of Labor Economics*, 4(2):243–256.

Basu, A. (2011). Estimating decision-relevant comparative effects using instrumental variables. *Statistics in Biosciences*, 3:6–27.

Basu, A. (2014). Estimating person-centered treatment (pet) effects using instrumental variables: an application to evaluating prostate cancer treatments. *Journal of Applied Econometrics*, 29(4):671–691.

Basu, A., Jones, A. M., and Dias, P. R. (2018). Heterogeneity in the impact of type of schooling on adult health and lifestyle. *Journal of Health Economics*, 57:1–14.

Basu, A. and Rathouz, P. J. (2005). Estimating marginal and incremental effects on health outcomes using flexible link and variance function models. *Biostatistics*, 6(1):93–109.

Baum, C. L. and Ford, W. F. (2004). The wage effects of obesity: a longitudinal study. *Journal of Health Economics*, 13(9):885–899.

Bazzoli, G. J. (1985). The early retirement decision: new empirical evidence on the influence of health. *Journal of Human Resources*, 20(2):214–234.

Benzeval, M. and Judge, K. (2001). Income and health: the time dimension. *Social Science and Medicine*, 52(9):1371–1390.

Bharadwaj, P., Pai, M. M., and Suziedelyte, A. (2017). Mental health stigma. *Economics Letters*, 159:57–60.

Blaine, B. (2008). Does depression cause obesity?: A meta-analysis of longitudinal studies of depression and weight control. *Journal of Health Psychology*, 13(8):1190–1197.

Blanchflower, D. G. and Oswald, A. J. (2004). Well-being over time in Britain and the USA. *Journal of Public Economics*, 88(7-8):1359–1386.

Bloom, D. E. and Canning, D. (2000). The health and wealth of nations. *Science*, 287(5456):1207–1209.

Bloom, D. E., Chatterji, S., Kowal, P., et al. (2015). Macroeconomic implications of population ageing and selected policy responses. *The Lancet*, 385(9968):649–657.

Böckerman, P., Cawley, J., Viinikainen, J., Lehtimäki, T., Rovio, S., Seppälä, I., Pehkonen, J., and Raitakari, O. (2019). The effect of weight on labor market outcomes: An application of genetic instrumental variables. *Health Economics*, 28(1):65–77.

Böckerman, P., Haapanen, M., and Jepsen, C. (2022). Dark passage: Mental health consequences of parental death. *Journal of Human Resources*, 58(1):196–229.

Bound, J., Brown, C., and Mathiowetz, N. (2001). Measurement error in survey data. In *Handbook of Econometrics*, volume 5, pages 3705–3843. Elsevier.

Bound, J., Schoenbaum, M., Stinebrickner, T. R., and Waidmann, T. (1999). The dynamic effects of health on the labor force transitions of older workers. *Labour Economics*, 6(2):179–202.

Brazenor, R. (2002). Disabilities and labour market earnings in australia. *Australian Journal of Labour Economics*, 5(3):319–334.

Brohan, E. and Thornicroft, G. (2010). Stigma and discrimination of mental health problems: workplace implications.

Brookes, A. J. (1999). The essence of snps. *Gene*, 234(2):177–186.

Bryan, M. L., Rice, N., Roberts, J., and Sechel, C. (2022). Mental health and employment: a bounding approach using panel data. *Oxford Bulletin of Economics and Statistics*, 84(5):1018–1051.

Bubonya, M., Cobb-Clark, D. A., and Ribar, D. C. (2019). The reciprocal relationship between depressive symptoms and employment status. *Economics & Human Biology*, 35:96–106.

Bünnings, C. and Tauchmann, H. (2015). Who opts out of the statutory health insurance? a discrete time hazard model for germany. *Health Economics*, 24(10):1331–1347.

Burrell, L. V., Mehlum, L., and Qin, P. (2022). Parental death by external causes during childhood and risk of psychiatric disorders in bereaved offspring. *Child and Adolescent Mental Health*, 27(2):122–130.

Busetta, G., Campolo, M. G., and Panarello, D. (2020). Weight-based discrimination in the italian labor market: an analysis of the interaction with gender and ethnicity. *The Journal of Economic Inequality*, 18:617–637.

Butler, J. S., Burkhauser, R. V., Mitchell, J. M., and Pincus, T. P. (1987). Measurement error in self-reported health variables. *The Review of Economics and Statistics*, 469(4):644–650.

Butterworth, P., Gill, S. C., Rodgers, B., Anstey, K. J., Villamil, E., and Melzer, D. (2006). Retirement and mental health: analysis of the australian national survey of mental health and well-being. *Social science & Medicine*, 62(5):1179–1191.

Bygren, M. (2004). Pay reference standards and pay satisfaction: what do workers evaluate their pay against? *Social Science Research*, 33(2):206–224.

Cai, J., Laporte, A., Zhang, L., et al. (2022). Impacts of absolute and relative income on self-rated health in urban and rural china. *International Journal of Health Services*, 52(1):129–140.

Cai, L. and Kalb, G. (2007). Health status and labour force status of older working-age australian men. *Australian Journal of Labour Economics*, 10(4):227–252.

Caliendo, M. and Gehrsitz, M. (2016). Obesity and the labor market: A fresh look at the weight penalty. *Economics & Human Biology*, 23:209–225.

Campbell, D. D., Green, M., Davies, N., Demou, E., Ward, J., Howe, L. D., Harrison, S., Johnston, K. J., Strawbridge, R. J., Popham, F., et al. (2021). Effects of increased body mass index on employment status: a mendelian randomisation study. *International Journal of Obesity*, 45(8):1790–1801.

Catania, L. S., Hetrick, S. E., Newman, L. K., and Purcell, R. (2011). Prevention and early intervention for mental health problems in 0–25 year olds: Are there evidence-based models of care? *Advances in Mental Health*, 10(1):6–19.

Cawley, J. (2004). The impact of obesity on wages. *Journal of Human Resources*, 39(2):451–474.

Cecchini, M. and Vuik, S. (2019). The heavy burden of obesity. *OECD Health Policy Studies*.

Centers for Disease Control and Prevention (2024). Obesity strategies: What can be done. https://www.cdc.gov/obesity/strategies/obesity-strategies-what-can-be-done.html. Accessed: 2025-04-29.

Chakravarty, S. (1990). *Social Ethical Index Numbers*. Springer-Verlag.

Chatterji, P., Alegría, M., Lu, M., and Takeuchi, D. (2007). Psychiatric disorders and labor market outcomes: evidence from the national latino and asian american study. *Health Economics*, 16(10):1069–1090.

Chatterji, P., Alegria, M., and Takeuchi, D. (2011). Psychiatric disorders and labor market outcomes: Evidence from the national comorbidity survey-replication. *Journal of Health Economics*, 30(5):858–868.

Clark, A. (2003). Inequality-aversion and income mobility: A direct test. *CNRS and DELTA-Federation Jourdan, Mimeo*, 4(12).

Clark, A. and Oswald, A. (1996). Satisfaction and comparison income. *Journal of Public Economics*, 61(3):359–381.

Cohen, S., Line, S., Manuck, S. B., Rabin, B. S., Heise, E. R., and Kaplan, J. R. (1997). Chronic social stress, social status, and susceptibility to upper respiratory infections in nonhuman primates. *Psychosomatic Medicine*, 59(3):213–221.

Conley, T. G., Hansen, C. B., and Rossi, P. E. (2012). Plausibly exogenous. *Review of Economics and Statistics*, 94(1):260–272.

Contoyannis, P., Jones, A. M., and Rice, N. (2004). The dynamics of health in the british household panel survey. *Journal of Applied Econometrics*, 19(4):473–503.

Cooperman, A., Sahgal, N., and Schiller, A. (2016). Israel's religiously divided society. *Pew Research Center*, 8.

Cui, X. and Chang, C.-T. (2021). How income influences health: decomposition based on absolute income and relative income effects. *International Journal of Environmental Research and Public Health*, 18(20):10738.

Dackehag, M., Gerdtham, U.-G., and Nordin, M. (2015). Productivity or discrimination? an economic analysis of excess-weight penalty in the swedish labor market. *The European Journal of Health Economics*, 16(6):589–601.

Dall, T. M., Sapra, T., Natale, Z., Livingston, T., and Chen, F. (2024). Assessing the economic impact of obesity and overweight on employers: identifying opportunities to improve work force health and well-being. *Nutrition & Diabetes*, 14(1):1–8.

Davey Smith, G. and Ebrahim, S. (2003). 'mendelian randomization': can genetic epidemiology contribute to understanding environmental determinants of disease? *International Journal of Epidemiology*, 32(1):1–22.

Davey Smith, G., Lawlor, D. A., Harbord, R., Timpson, N., Day, I., and Ebrahim, S. (2007). Clustered environments and randomized genes: a fundamental distinction between conventional and genetic epidemiology. *PLoS Medicine*, 4(12):e352.

Deaton, A. (2001). *Relative deprivation, inequality, and mortality. Working Paper*. 275. Princeton University, Woodrow Wilson School of Public and International Affairs, Center for Health and Wellbeing.

Department for Work and Pensions and Department of Health (2016). Improving lives: The work, health and disability green paper. https://www.gov.uk/government/consultations/work-health-and-disability-improving-lives. Accessed: 2025-04-29.

Destri, K., Alves, J., Gregório, M. J., Dias, S. S., Henriques, A. R., Mendonça, N., Canhão, H., and Rodrigues, A. M. (2022). Obesity-attributable costs of absenteeism among working adults in portugal. *BMC Public Health*, 22(1):978.

Diabetes UK (2024). Unemployed could receive weight-loss jabs to return to work. Accessed: 2025-05-01.

Disney, R., Emmerson, C., and Wakefield, M. (2006). Ill health and retirement in britain: A panel data-based analysis. *Journal of Health Economics*, 25(4):621–649.

Drewnowski, A. and Barratt-Fornell, A. (2004). Do healthier diets cost more? *Nutrition Today*, 39(4):161–168.

Duesenberry, J. (1949). *Income, saving, and the theory of consumer behavior.* Harvard Economic Studies ; v. 87. Harvard University Press, Cambridge.

Ecob, R. and Smith, G. D. (1999). Income and health: what is the nature of the relationship? *Social Science and Medicine*, 48(5):693–705.

Eibner, C. and Evans, W. N. (2005). Relative deprivation, poor health habits, and mortality. *Journal of Human Resources*, 40(3):591–620.

Ellaway, A., Benzeval, M., Green, M., Leyland, A., and Macintyre, S. (2012). Getting sicker quicker": does living in a more deprived neighbourhood mean your health deteriorates faster? *Health and Place*, 18(2):132–137.

Ettner, S. L. (1996). New evidence on the relationship between income and health. *Journal of Health Economics*, 15(1):67–85.

Ettner, S. L., Frank, R. G., and Kessler, R. C. (1997). The impact of psychiatric disorders on labor market outcomes. *ILR Review*, 51(1):64–81.

Fadare, O., Zanello, G., and Srinivasan, C. (2023). Stressor or succour? examining the association between conflict, livestock assets, and farmers' mental health in nigeria. *Economics & Human Biology*, 49:101234.

Feinstein, J. S. (1993). The relationship between socioeconomic status and health: A review of the literature. *The Milbank Quarterly*, 71(2):279–322.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7(2):117–140.

Fiscella, K. and Franks, P. (1997). Poverty or income inequality as predictor of mortality: longitudinal cohort study. *British Medical Journal*, 314(7096):1724.

Frank, R. H. (1985). Choosing the right pond: Human behavior and the quest for status. *Public Choice*, 56(1):84–88.

Frijters, P., Haisken-DeNew, J. P., and Shields, M. A. (2005). The causal effect of income on health: Evidence from german reunification. *Journal of Health Economics*, 24(5):997–1017.

Frijters, P., Johnston, D., and Shields, M. (2010). Mental health and labour market participation: Evidence from iv panel data models. *Institute for the Study of Labor (IZA), IZA Discussion Papers*.

Fuller, R. C. (1996). *Religion and Wine: A Cultural History of Wine Drinking in the United States*. University of Tennessee Press.

García-Gómez, P., Jones, A. M., and Rice, N. (2010). Health effects on labour market exits and entries. *Labour Economics*, 17(1):62–76.

Garrido, M. M., Deb, P., Burgess Jr, J. F., and Penrod, J. D. (2012). Choosing models for health care cost analyses: issues of nonlinearity and endogeneity. *Health Services Research*, 47(6):2377–2397.

Gibson, M., Hearty, W., and Craig, P. (2020). The public health effects of interventions similar to basic income: A scoping review. *The Lancet Public Health*, 5(3):e165–e176.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119.

Goulão, C., Lacomba, J. A., Lagos, F., and Rooth, D.-O. (2024). Weight, attractiveness, and gender when hiring: A field experiment in spain. *Journal of Economic Behavior & Organization*, 218:132–145.

Gravelle, H. and Sutton, M. (2009). Income, relative income, and self-reported health in britain 1979–2000. *Journal of Health Economics*, 18(2):125–145.

Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political Economy*, 80(2):223–255.

Hamilton, V. H., Merrigan, P., and Dufresne, É. (1997). Down and out: estimating the relationship between mental health and unemployment. *Health Economics*, 6(4):397–406.

Han, E., Norton, E. C., and Powell, L. M. (2011). Direct and indirect effects of body weight on adult wages. *Economics & Human Biology*, 9(4):381–392.

Hausman, J. A. and Taylor, W. E. (1981). Panel data and unobservable individual effects. *Econometrica: Journal of the Econometric Society*, 49(6):1377–1398.

Heckman, J. and Singer, B. (1984). A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica: Journal of the Econometric Society*, 52(2):271–320.

Heckman, J. J. and Robb Jr, R. (1985). Alternative methods for evaluating the impact of interventions: An overview. *Journal of Econometrics*, 30(1-2):239–267.

Heckman, J. J. and Vytlacil, E. J. (1999). Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of The National Academy of Sciences*, 96(8):4730–4734.

Hey, J. D. and Lambert, P. J. (1980). Relative deprivation and the gini coefficient: comment. *The Quarterly Journal of Economics*, 95(3):567–573.

Holt, C. L., Caplan, L., Schulz, E., Blake, V., Southward, P., Buckner, A., and Lawrence, H. (2009). Role of religion in cancer coping among african americans: A qualitative examination. *Journal of Psychosocial Oncology*, 27(2):248–273.

Idler, E. L., Musick, M. A., Ellison, C. G., George, L. K., Krause, N., Ory, M. G., Pargament, K. I., Powell, L. H., Underwood, L. G., and Williams, D. R. (2003). Measuring multiple dimensions of religion and spirituality for health research: Conceptual background and findings from the 1998 general social survey. *Research on Aging*, 25(4):327–365.

Imbens, G. and Rubin, D. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, Cambridge.

IPPR (2025). Children with poor mental health are two-thirds more likely to have reduced ability to work as adults. Accessed: 2025-05-23.

Jenkins, S. P. (1995). Easy estimation methods for discrete-time duration models. *Oxford Bulletin of Economics and Statistics*, 57(1):129–138.

Jenkins, S. P. (1997). Estimation of discrete time (grouped duration data) proportional hazards models: pgmhaz. *Stata Technical Bulletin Reprints, STB*, 17(7):109–121.

Jones, A. M. (2009). Panel data methods and applications to health economics. *Palgrave Handbook of Econometrics: Applied Econometrics*, 2:557–631.

Jones, A. M. and Wildman, J. (2008). Health, income and relative deprivation: Evidence from the bhps. *Journal of Health Economics*, 27(2):308–324.

Kanazawa, S. (2013). Childhood intelligence and adult obesity. *Obesity*, 21(3):434–440.

Kanazawa, S. (2014). Intelligence and obesity: which way does the causal direction go? *Current Opinion in Endocrinology, Diabetes, and Obesity*, 21(5):339–344.

Kaplan, G. A., Pamuk, E. R., Lynch, J. W., Cohen, R. D., and Balfour, J. L. (1996). Inequality in income and mortality in the united states: analysis of mortality and potential pathways. *British Medical Journal*, 312(7037):999–1003.

Kawachi, I., Levine, S., Miller, S. M., Lasch, K., and Amick III, B. C. (1994). *Income Inequality and Life Expectancy: Theory, Research, and Policy*. The Health Institute, New England Medical Center, Boston.

Kennedy, B. P., Kawachi, I., and Prothrow-Stith, D. (1996). Income distribution and mortality: cross sectional ecological study of the robin hood index in the united states. *British Medical Journal*, 312(7037):1004–1007.

Kesaite, V. and Greve, J. (2024). The impact of excess body weight on employment outcomes: A systematic review of the evidence. *Economics & Human Biology*, 54:101398.

Kessler, R. C., Andrews, G., Colpe, L. J., Hiripi, E., Mroczek, D. K., Normand, S.-L., Walters, E. E., and Zaslavsky, A. M. (2002). Short screening scales to monitor population prevalences and trends in non-specific psychological distress. *Psychological Medicine*, 32(6):959–976.

Kessler, R. C., Mickelson, K. D., and Williams, D. R. (1999). The prevalence, distribution, and mental health correlates of perceived discrimination in the united states. *Journal of Health and Social Behavior*, 40(3):208–230.

Kim, D. and Leigh, J. P. (2010). Estimating the effects of wages on obesity. *Journal of Occupational and Environmental Medicine*, 52(5):495–500.

Koenig, H. G. (2012). Religion, spirituality, and health: The research and clinical implications. *International Scholarly Research Notices*.

Kühn, J. and Wolbring, T. (2024). Beauty pays, but not under all circumstances: Evidence on gendered hiring discrimination from a novel experimental treatment using deepfakes. *Research in Social Stratification and Mobility*, 94:100992.

Lancaster, T. (1990). *The Econometric Analysis of Transition Data*. 17. Cambridge University Press.

Lawlor, D. A., Harbord, R. M., Sterne, J. A., Timpson, N., and Davey Smith, G. (2008a). Mendelian randomization: using genes as instruments for making causal inferences in epidemiology. *Statistics in Medicine*, 27(8):1133–1163.

Lawlor, D. A., Windmeijer, F., and Davey Smith, G. (2008b). Is mendelian randomization 'lost in translation?': Comments on 'mendelian randomization equals instrumental variable analysis with genetic instruments' by wehby et al. *Statistics in Medicine*, 27(15):2750–2755.

Lee, D. S., McCrary, J., Moreira, M. J., and Porter, J. (2022). Valid t-ratio inference for iv. *American Economic Review*, 112(10):3260–3290.

Lee, S. and Lee, W. (2025). A review of mendelian randomization: Assumptions, methods, and application to obesity-related diseases. *Journal of Obesity & Metabolic Syndrome*, 34(1):14.

Lee, Y.-C., Chatterton, M. L., Magnus, A., Mohebbi, M., Le, L. K.-D., and Mihalopoulos, C. (2017). Cost of high prevalence mental disorders: Findings from the 2007 australian national survey of mental health and wellbeing. *Australian & New Zealand Journal of Psychiatry*, 51(12):1198–1211.

Lemeshow, S., May, S., and Hosmer Jr, D. W. (2011). *Applied Survival Analysis: Regression Modeling of Time-to-event Data*. John Wiley & Sons.

Levin, J. (2013). Religious behavior, health, and well-being among israeli jews: Findings from the european social survey. *Psychology of Religion and Spirituality*, 5(4):272.

Li, P., Chen, X., Stafford, F., and Ou, J. (2022). Body shape and stable employment opportunity analysis of china's nonagricultural labor market. *SSM-Population Health*, 17:101014.

Liddiard, P. (1978). *An Ageing Population: a Reader and Sourcebook*. Hodder & Stoughton.

Lindeboom, M., Lundborg, P., and Van Der Klaauw, B. (2010). Assessing the impact of obesity on labor market outcomes. *Economics & Human Biology*, 8(3):309–319.

Lindley, J. and Lorgelly, P. (2005). The relative income hypothesis: does it exist over time? evidence from the bhps. Working Paper ISSN 1749-8368, Department of Economics, University of Sheffield.

Locke, A. E., Kahali, B., Berndt, S. I., Justice, A. E., Pers, T. H., Day, F. R., Powell, C., Vedantam, S., Buchkovich, M. L., Yang, J., et al. (2015). Genetic studies of body mass index yield new insights for obesity biology. *Nature*, 518(7538):197–206.

Lu, C., Frank, R. G., Liu, Y., Shen, J., et al. (2009). The impact of mental health on labour market outcomes in china. *Journal of Mental Health Policy and Economics*, 12(3):157.

Lundborg, P., Bolin, K., Höjgård, S., and Lindgren, B. (2006). Obesity and occupational attainment among the 50+ of europe. In *The Economics of Obesity*, pages 219–251. Emerald Group Publishing Limited.

Lynch, J. W. and Kaplan, G. A. (1997). Understanding how inequality in the distribution of income affects health. *Journal of Health Psychology*, 2(3):297–314.

Macleod, J. and Smith, G. D. (2003). Psychosocial factors and public health: a suitable case for treatment? *Journal of Epidemiology & Community Health*, 57(8):565–570.

Mangyo, E. and Park, A. (2011). Relative deprivation and health which reference groups matter? *Journal of Human Resources*, 46(3):459–481.

Marmot, M. G., Shipley, M. J., and Rose, G. (1984). Inequalities in death—specific explanations of a general pattern? *The Lancet*, 323(8384):1003–1006.

McBride, M. (2001). Relative-income effects on subjective well-being in the cross-section. *Journal of Economic Behavior and Organization*, 45(3):251–278.

McDaid, D. and Park, A.-L. (2022). Understanding the economic value and impacts on informal carers of people living with mental health conditions. *International Journal of Environmental Research and Public Health*, 19(5):2858.

McManus, S., Bebbington, P., Jenkins, R., and Brugha, T. (2016). *Mental health and wellbeing in England: Adult Psychiatric Morbidity Survey 2014*. NHS Digital, Leeds.

Melzer, D., Buxton, J., and Villamil, E. (2004). Decline in common mental disorder prevalence in men during the sixth decade of life: evidence from the national psychiatric morbidity survey. *Social Psychiatry and Psychiatric Epidemiology*, 39(1):33–38.

Mery, G., Wodchis, W. P., and Laporte, A. (2016). The determinants of the propensity to receive publicly funded home care services for the elderly in canada: a panel two-stage residual inclusion approach. *Health Economics Review*, 6(1):1–18.

Mitra, S., Gao, Q., Chen, W., et al. (2020). Health, work, and income among middle-aged and older adults: A panel analysis for china. *The Journal of the Economics of Ageing*, 17:100255.

Moler-Zapata, S., Grieve, R., Basu, A., and O'Neill, S. (2023). How does a local instrumental variable method perform across settings with instruments of differing strengths? a simulation study and an evaluation of emergency surgery. *Health Economics*, 32(9):2113–2126.

Morris, S. (2007). The impact of obesity on employment. *Labour Economics*, 14(3):413–433.

Mossakowski, K. N. (2009). The influence of past unemployment duration on symptoms of depression among young women and men in the united states. *American Journal of Public Health*, 99(10):1826–1832.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica: Journal of The Econometric Society*, 46(1):69–85.

Neil, A. L., Carr, V. J., Mihalopoulos, C., Mackinnon, A., and Morgan, V. A. (2014). Costs of psychosis in 2010: Findings from the second australian national survey of psychosis. *Australian & New Zealand Journal of Psychiatry*, 48(2):169–182.

Nerlove, M. (2005). *Essays in Panel Data Econometrics*. Cambridge University Press.

Newhouse, J. P. and McClellan, M. (1998). Econometrics in outcomes research: the use of instrumental variables. *Annual Review of Public Health*, 19(1):17–34.

Ngui, E. M., Khasakhala, L., Ndetei, D., and Roberts, L. W. (2010). Mental disorders, health inequalities and ethics: A global perspective. *International Review of Psychiatry*, 22(3):235–244.

NHLBI (2024). Assessing your weight and health risk. Accessed: March 25, 2025.

NHS (2023). Obesity.

NHS Digital (2009). Health survey for england 2008: Physical activity and fitness. Accessed: March 25, 2025.

NHS England (2024). England's nhs mental health services treat record 3.8 million people last year. Accessed: 01 May 2025.

NICE (2014). Obesity: identification, assessment and management.

NICE (2025). Overweight and obesity management. Accessed: 2025-04-22.

Nicoletti, C. and Rondinelli, C. (2010). The (mis) specification of discrete duration models with unobserved heterogeneity: a monte carlo study. *Journal of Econometrics*, 159(1):1–13.

Norton, E. C. and Han, E. (2008). Genetic information, obesity, and labor market outcomes. *Journal of Health Economics*, 17(9):1089–1104.

Olesen, S. C., Butterworth, P., Leach, L. S., Kelaher, M., and Pirkis, J. (2013). Mental health affects future employment as job loss affects mental health: findings from a longitudinal population study. *BMC Psychiatry*, 13:1–9.

Palmer, T. M., Lawlor, D. A., Harbord, R. M., Sheehan, N. A., Tobias, J. H., Timpson, N. J., Smith, G. D., and Sterne, J. A. (2012). Using multiple genetic variants as instrumental variables for modifiable risk factors. *Statistical Methods in Medical Research*, 21(3):223–242.

Parker, M., Bucknall, M., Jagger, C., and Wilkie, R. (2020). Population-based estimates of healthy working life expectancy in england at age 50 years: analysis of data from the english longitudinal study of ageing. *The Lancet Public Health*, 5(7):e395–e403.

Popkin, B. M., Du, S., Green, W. D., Beck, M. A., Algaith, T., Herbst, C. H., Alsukait, R. F., Alluhidan, M., Alazemi, N., and Shekar, M. (2020). Individuals with obesity and covid-19: a global perspective on the epidemiology and biological relationships. *Obesity Reviews*, 21(11):e13128.

Porru, F., Burdorf, A., and Robroek, S. J. (2019). The impact of depressive symptoms on exit from paid employment in europe: a longitudinal study with 4 years follow-up. *European Journal of Public Health*, 29(1):134–139.

Preston, S. H. (1975). The changing relation between mortality and level of economic development. *Population Studies*, 29(2):231–248.

Productivity Commission (2020). Mental health, productivity commission inquiry report. Accessed: 23 May 2025.

Reche, E., König, H.-H., and Hajek, A. (2019). Income, self-rated health, and morbidity: A systematic review of longitudinal studies. *International Journal of Environmental Research and Public Health*, 16(16):2884.

Robert, S. A. (1998). Community-level socioeconomic status effects on adult health. *Journal of Health and Social Behavior*, 39(1):18–37.

Robinson, P. M. (1988). Root-n-consistent semiparametric regression. *Econometrica: Journal of the Econometric Society*, 56(4):931–954.

Ruckert, A., Huynh, C., and Labonté, R. (2017). Reducing health inequities: Is universal basic income the way forward? *J. Public Health*, 40(1):3–7.

Runciman, W. G. (1966). *Relative deprivation and social justice: A study of attitudes to social inequality in twentieth-century England*, volume 13. Routledge/Thoemms Press.

Rüsch, N., Angermeyer, M. C., and Corrigan, P. W. (2005). Mental illness stigma: Concepts, consequences, and initiatives to reduce stigma. *European Psychiatry*, 20(8):529–539.

Sainsbury, R., Irvine, A., Aston, J., Wilson, S., Williams, C., and Sinclair, A. (2008). Mental health and employment. *White Rose*.

Sanderson, E., Glymour, M. M., Holmes, M. V., Kang, H., Morrison, J., Munafò, M. R., Palmer, T., Schooling, C. M., Wallace, C., Zhao, Q., et al. (2022). Mendelian randomization. *Nature Reviews Methods Primers*, 2(1):6.

Senik, C. (2008). Ambition and jealousy: Income interactions in the 'old' europe versus the 'new'europe and the united states. *Economica*, 75(299):495–513.

Seral-Cortes, M., Sabroso-Lasa, S., De Miguel-Etayo, P., Gonzalez-Gross, M., Gesteiro, E., Molina-Hidalgo, C., De Henauw, S., Gottrand, F., Mavrogianni, C., Manios, Y., et al. (2021). Development of a genetic risk score to predict the risk of overweight and obesity in european adolescents from the helena study. *Scientific Reports*, 11(1):3067.

Siddiqui, S. and Ali Shah, S. A. (1997). Siddiqui-shah depression scale (ssds): development and validation. *Psychology and Developing Societies*, 9(2):245–262.

Smith, G. D. (2011). Random allocation in observational data: how small but robust effects could facilitate hypothesis-free causal inference. *Epidemiology*, 22(4):460–463.

Smith, J. P. (1999). Healthy bodies and thick wallets: The dual relationship between health and economic status. *The Journal of Economic Perspectives*, 13(2):145.

Speliotes, E. K., Willer, C. J., Berndt, S. I., Monda, K. L., Thorleifsson, G., Jackson, A. U., Allen, H. L., Lindgren, C. M., Luan, J., Mägi, R., et al. (2010). Association analyses of 249,796 individuals reveal 18 new loci associated with body mass index. *Nature Genetics*, 42(11):937–948.

Staiger, D. and Stock, J. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557–586.

Steele, F., French, R., and Bartley, M. (2013). Adjusting for selection bias in longitudinal analyses using simultaneous equations modeling: the relationship between employment transitions and mental health. *Epidemiology*, 24(5):703–711.

Stern, S. (1989). Measuring the effect of disability on labor force participation. *Journal of Human Resources*, 24(3):361–395.

Stock, J. H. and Wise, D. A. (1988). Pensions, the option value of work, and retirement.

Subramanian, S. V. and Kawachi, I. (2004). Income inequality and health: what have we learned so far? *Epidemiologic Reviews*, 26(1):78–91.

Tchetgen, E. J. T., Walter, S., Vansteelandt, S., Martinussen, T., and Glymour, M. (2015). Instrumental variable estimation in a survival context. *Epidemiology (Cambridge, Mass.)*, 26(3):402.

Terza, J. V. (2018). Two-stage residual inclusion estimation in health services research and health economics. *Health Services Research*, 53(3):1890–1899.

Terza, J. V., Basu, A., and Rathouz, P. J. (2008). Two-stage residual inclusion estimation: addressing endogeneity in health econometric modeling. *Journal of Health Economics*, 27(3):531–543.

Thompson, D. J., Wells, D., Selzam, S., Peneva, I., Moore, R., Sharp, K., Tarran, W. A., Beard, E. J., Riveros-Mckay, F., Giner-Delgado, C., Palmer, D., Seth, P., Harrison, J., Futema, M., Consortium, G. E. R., McVean, G., Plagnol, V., Donnelly, P., and Weale, M. E. (2022). Uk biobank release and systematic evaluation of optimised polygenic risk scores for 53 diseases and quantitative traits. *MedRxiv*.

Tompa, E. (2002). *The impact of health on productivity: empirical evidence and policy implications*. CIGI.

Tyrrell, J., Richmond, R. C., Palmer, T. M., Feenstra, B., Rangarajan, J., Metrustry, S., Cavadino, A., Paternoster, L., Armstrong, L. L., De Silva, N. M. G., et al. (2016). Genetic evidence for causal relationships between maternal obesity-related traits and birth weight. *Jama*, 315(11):1129–1140.

Uffelmann, E., Huang, Q., and Munung, N. e. a. (2021). Genome-wide association studies. *Nat Rev Methods Primers*, 1:59.

Veenstra, G. and Vanzella-Yang, A. (2020). Family income and self-rated health in canada: Using fixed effects models to control for unobserved confounders and investigate causal temporality. *Social Science and Medicine*, 250:112884.

Virtanen, M., Jokela, M., Lallukka, T., and et al. (2020). Long working hours and change in body weight: analysis of individual-participant data from 19 cohort studies. *International Journal of Obesity*, 44:1368–1375.

Von Hinke, S., Smith, G. D., Lawlor, D. A., Propper, C., and Windmeijer, F. (2016). Genetic markers as instrumental variables. *Journal of Health Economics*, 45:131–148.

von Hinke, Kessler Scholder, S., Wehby, G. L., Lewis, S., and Zuccolo, L. (2014). Alcohol exposure in utero and child academic achievement. *The Economic Journal*, 124(576):634–667.

Vásquez, F., Paraje, G., and Estay, M. (2013). Income-related inequality in health and health care utilization in chile, 2000-2009. *Revista Panamericana de Salud Pública*, 33(2):98–106.

Wan, F., Small, D., Bekelman, J. E., and Mitra, N. (2015). Bias in estimating the causal hazard ratio when using two-stage instrumental variable methods. *Statistics in Medicine*, 34(14):2235–2265.

Ware, J. E., Snow, K. K., Kosinski, M., and Gandek, B. (1993). Sf-36 health survey. *Manual and Interpretation Guide*, 2(6):473–83.

WHO (2024). *Operational framework for monitoring social determinants of health equity*. World Health Organization.

Wilkins, R. (2004). The effects of disability on labour force status in australia. *Australian Economic Review*, 37(4):359–382.

Wilkinson, R. (1996). *Unhealthy Societies: The Afflictions of Inequality (1st ed.)*. Routledge.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

World Health Organisation (2021). *World Report on Disability 2021*. World Health Organisation.

World Health Organisation (2022). Ageing and health. https://www.who.int/news-room/fact-sheets/detail/ageing-and-health.

Yengo, L., Sidorenko, J., Kemper, K. E., Zheng, Z., Wood, A. R., Weedon, M. N., Frayling, T. M., Hirschhorn, J., Yang, J., Visscher, P. M., et al. (2018). Meta-analysis of genome-wide association studies for height and body mass index in 700000 individuals of european ancestry. *Human Molecular Genetics*, 27(20):3641–3649.

Ying, A., Xu, R., and Murphy, J. (2019). Two-stage residual inclusion for survival data and competing risks—an instrumental variable approach with application to seer-medicare linked data. *Statistics in Medicine*, 38(10):1775–1801.

Zhang, X., Zhao, X., and Harris, A. (2009). Chronic diseases and labour force participation in australia. *Journal of Health Economics*, 28(1):91–108.

Zucchelli, E., Harris, A., Jones, A. M., and Rice, N. (2010). The effects of health shocks on labour market exits: Evidence from the hilda survey. *Australian Journal of Labour Economics*, 13(2):191–218.

Appendix A

Appendix to Chapter 1

A.1 Full-model: LPM Regression

Table A.1 LPM Regression: Effect of Mental Health on Early Retirement

VARIABLES	**Retired**
Lagged mental health	-0.0017***
	(0.0001)
Income	-0.0005***
	(0.0001)
Female	0.0360***
	(0.0041)
Married	0.0398***
	(0.0048)
Degree	-0.0281***
	(0.0049)
Household size	-0.0258***
	(0.0020)
Local unemployment rate	-0.0080***
	(0.0020)
Age 53	-0.1055***
	(0.0088)
Age 54	-0.1035***

Continued on next page

Table A.1 – Continued

VARIABLES	Retired
	(0.0089)
Age 55	-0.0914***
	(0.0090)
Age 56	-0.0567***
	(0.0092)
Age 57	-0.0333***
	(0.0093)
Age 58	-0.0184*
	(0.0094)
Age 59	0.0119
	(0.0095)
Age 60	0.0583***
	(0.0097)
Age 61	0.1052***
	(0.0098)
Age 62	0.1388***
	(0.0099)
Age 63	0.1757***
	(0.0106)
Age 64	0.2228***
	(0.0114)
Constant	0.4176***
	(0.0159)
Observations	32,057
R-squared	0.1067

Standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

A.2 Full-model: 2SLS Regression

Table A.2 2SLS Regression: Effect of Mental Health on Early Retirement

VARIABLES	Retired
Lagged mental health	0.0036
	(0.0032)
Degree	-0.0394***
	(0.0073)
Income	-0.0007***
	(0.0001)
Household size	-0.0172***
	(0.0030)
Married	0.0088
	(0.0169)
Local unemployment rate	0.0008
	(0.0024)
Age 52	-0.4115***
	(0.0167)
Age 53	-0.3960***
	(0.0173)
Age 54	-0.3924***
	(0.0175)
Age 55	-0.3820***
	(0.0162)
Age 56	-0.3422***
	(0.0172)
Age 57	-0.3213***
	(0.0159)
Age 58	-0.3061***
	(0.0151)
Age 59	-0.2735***
	(0.0155)

Continued on next page

Table A.2 – Continued

VARIABLES	Retired
Age 60	-0.2255***
	(0.0152)
Age 61	-0.1737***
	(0.0152)
Age 62	-0.1354***
	(0.0146)
Age 63	-0.1178***
	(0.0148)
Age 64	-0.0687***
	(0.0157)
Constant	0.3077
	(0.2662)
Observations	29,432
R-squared (Centered)	0.0870
R-squared (Uncentered)	0.2659
Root MSE	0.3792
Underidentification test (LM stat.)	38.794
Underidentification p-value	0.0000
Weak identification test (CD F stat.)	38.819

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

A.3 Full-model: Complementary log-log model

Table A.3 Complementary Log-Log Regression: Effect of Mental Health on Early Retirement

VARIABLES	Early Retirement (Hazard)
Lagged mental health	-0.0083***
	(0.0028)
Initial mental health	-0.0076***

Continued on next page

Table A.3 – Continued

VARIABLES	**Early Retirement (Hazard)**
	(0.0029)
Income	0.0006
	(0.0005)
Female	0.0620
	(0.0836)
Married	0.0613
	(0.0989)
Degree	0.5969***
	(0.1109)
Professional/Managerial Jobs	-1.9837***
	(0.1484)
Negative Health Shock	-0.2812
	(0.1941)
Household Size	-0.1893***
	(0.0544)
Local unemployment rate	-0.0112
	(0.0390)
Age 53	-1.5787***
	(0.4287)
Age 54	-1.7052***
	(0.3763)
Age 55	-1.4337***
	(0.2962)
Age 56	-0.4093**
	(0.1926)
Age 57	-0.6121***
	(0.1992)
Age 58	-0.6215***
	(0.1981)
Age 59	-0.5594***
	(0.1898)
Age 60	-0.3353*

Continued on next page

Table A.3 – Continued

VARIABLES	Early Retirement (Hazard)
	(0.1769)
Age 61	-0.0721
	(0.1680)
Age 62	-0.1441
	(0.1715)
Age 63	-0.1332
	(0.1805)
Age 64	-0.0783
	(0.1857)
Constant	-0.3713
	(0.3381)
Observations	9,054

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

A.4 Full-model: Second-Stage of 2SRI

Table A.4 The second-stage of the Two-Stage Residual Inclusion

VARIABLES	Early Retirement
Lagged mental health	-0.011***
	(0.003)
Initial mental health	-0.008***
	(0.003)
Income	-0.001
	(0.001)
Married	0.117
	(0.100)
Degree	-0.199*
	(0.105)

Continued on next page

Table A.4 – Continued

VARIABLES	Early Retirement
Negative health shock	-0.175
	(0.194)
Household size	-0.138**
	(0.055)
Local unemployment rate	0.149***
	(0.041)
age53	-1.331***
	(0.429)
age54	-1.589***
	(0.376)
age55	-1.471***
	(0.296)
age56	-0.480**
	(0.194)
age57	-0.732***
	(0.201)
age58	-0.776***
	(0.200)
age59	-0.724***
	(0.191)
age60	-0.505***
	(0.177)
age61	-0.234
	(0.168)
age62	-0.304*
	(0.172)
age63	-0.417**
	(0.181)
age64	-0.319*
	(0.185)
Xuhat	0.052
	(0.059)

Continued on next page

Table A.4 – Continued

VARIABLES	Early Retirement
Constant	-4.847
	(4.558)
Observations	7,335

Standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

A.5 Full-model: Sub-sample of females and males

Table A.5 LPM estimates of early retirement on mental health and covariates, by gender

	Female	Male
Lagged mental health	-0.001***	-0.002***
	(0.001)	(0.001)
Degree	-0.031***	-0.012*
	(0.007)	(0.007)
Income	-0.001***	-0.001***
	(0.001)	(0.001)
Household size	-0.012***	-0.016***
	(0.003)	(0.003)
Married	0.067***	-0.023***
	(0.007)	(0.007)
Local unemployment rate	-0.012***	-0.007**
	(0.003)	(0.003)
Negative health shock	-0.032**	-0.008
	(0.014)	(0.013)
Age 53 to 64 Controls	Included	Included
Observations	16,192	15,865
R-squared	0.131	0.152

Notes: Robust standard errors in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.6 2SLS Estimates of Mental Health on Retirement, Instrumented by Death of Close Friend

	Dependent variable: Retired	
Variable	Female (N=14,926)	Male (N=14,506)
Lagged mental health	0.0116	-0.0004
	(0.0090)	(0.0034)
Degree	-0.0635**	-0.0185**
	(0.0210)	(0.0073)
Income	-0.0009***	-0.0006***
	(0.0002)	(0.0001)
Household size	-0.0040	-0.0174***
	(0.0082)	(0.0033)
Married	0.0100	-0.0241*
	(0.0469)	(0.0140)
Local unemployment rate	-0.0013	0.0024
	(0.0041)	(0.0030)
Negative health shock	0.0384	0.0069
	(0.0488)	(0.0214)
Age 52 to 64 Controls	Included	Included
Constant	-0.2623	0.6138*
	(0.6639)	(0.2568)
Number of observations	14,926	14,506
Underidentification test (Anderson LM)	7.91	39.87
Weak identification (Cragg-Donald F)	7.90	39.92

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, . $p < 0.10$

Table A.7 Complementary Log-Log Regression on Retirement by Sex

	Dependent variable: Early retirement (_d)	
Variable	Female (N=4,126)	Male (N=4,928)
Lagged mental health	-0.0057	-0.0147***
	(0.0041)	(0.0041)
Initial mental health	-0.0068	-0.0080*
	(0.0042)	(0.0041)
Income	-0.0006	-0.0010
	(0.0008)	(0.0008)
Married	0.2223	-0.1042
	(0.1428)	(0.1390)
Degree	-0.3473*	0.0174
	(0.1651)	(0.1380)
Negative health shock	-0.5775*	-0.0095
	(0.3400)	(0.2378)
Household size	-0.0259	-0.2362***
	(0.0843)	(0.0734)
Local unemployment rate	0.0622	-0.0477
	(0.0578)	(0.0537)
Age 53	-0.5912	-2.8383**
Age 54	-0.9307*	-2.2630***
Age 55	-0.7205	-1.8665***
Age 56	0.1301	-0.5563**
Age 57	0.2913	-1.1537***
Age 58	-0.0018	-0.8705***
Age 59	0.1140	-0.7936***
Age 60	0.5169	-0.7007***
Age 61	0.7893*	-0.4382**
Age 62	0.7452*	-0.5104**
Age 63	0.4823	-0.2896
Age 64	0.7171*	-0.2862
Constant	-2.2369***	0.6060
	(0.5493)	(0.4526)
Number of obs	4,126	4,928
Log likelihood	-1020.27	-1131.84
LR chi2(20)	79.18	136.95
Prob > chi2	0.0000	0.0000

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, . $p < 0.10$

Table A.8 2SRI Complementary Log-Log Regression on Retirement by Sex

	Female (N=3,320)	**Male** (N=4,015)
Lagged mental health	-0.0063	-0.0164***
	(0.0038)	(0.0038)
Residual from first stage (Xuhat)	0.0072	0.0688
	(0.1205)	(0.0670)
Initial mental health	-0.0080**	-0.0094**
	(0.0040)	(0.0038)
Income	-0.0006	-0.0012
	(0.0008)	(0.0009)
Married	0.2147	-0.1149
	(0.1438)	(0.1395)
Degree	-0.3685**	-0.0101
	(0.1651)	(0.1383)
Negative health shock	-0.4792	0.0398
	(0.3403)	(0.2381)
Household size	0.0090	-0.2078***
	(0.0857)	(0.0733)
Local unemployment rate	0.2065***	0.1027*
	(0.0603)	(0.0558)
Age 53	-0.4963	-2.4146**
Age 54	-1.0005*	-1.9276***
Age 55	-0.9464**	-1.7801***
Age 56	-0.0708	-0.5921**
Age 57	-0.0556	-1.1802***
Age 58	-0.3461	-0.8813***
Age 59	-0.2607	-0.8656***
Age 60	0.1000	-0.7887***
Age 61	0.3809	-0.5138**
Age 62	0.3692	-0.6467***
Age 63	-0.0481	-0.4738**
Age 64	0.1268	-0.4338**
Constant	-2.9029	-5.1366
	(9.1336)	(5.3089)
Log Likelihood	-947.359	-1059.165
Observations	3,320	4,015

Robust standard errors in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Appendix B

Appendix to Chapter 2

B.1 Full-model: Absolute and Relative Income Hypotheses - Ordered Probit

Table B.1 Absolute and Relative Income Hypotheses - Ordered Probit

	Absolute Income	**Relative Income**
loglincome	0.008	0.002
	(0.015)	(0.016)
dep	–	-0.102***
	–	(0.024)
hebrew	0.276***	0.258***
	(0.073)	(0.074)
arabic	-0.254***	-0.232***
	(0.086)	(0.087)
english	-0.813***	-0.821***
	(0.080)	(0.080)
female	-0.055*	-0.061**
	(0.029)	(0.029)
academicdegree	0.380***	0.362***
	(0.038)	(0.039)
married	0.032	0.037

Continued on next page

	Absolute Income	**Relative Income**
	(0.039)	(0.039)
age2129	1.902***	1.934***
	(0.064)	(0.065)
age3039	1.361***	1.372***
	(0.048)	(0.049)
age4049	0.951***	0.955***
	(0.045)	(0.046)
age5059	0.443***	0.446***
	(0.041)	(0.042)
meanloglincome	0.198***	0.136***
	(0.024)	(0.035)
meandep	–	-0.031
	–	(0.050)
/cut1	-0.551	-1.396
	(0.224)	(0.363)
/cut2	1.159	0.345
	(0.222)	(0.362)
/cut3	3.005	2.207
	(0.223)	(0.362)
$/\sigma_u^2$	1.211	1.229
	(0.050)	(0.051)
Observations	20,785	19,063
Number of ID	5,773	5,368

Standard errors in parentheses

**** p<0.01, ** p<0.05, * p<0.1*

B.2 Sub-sample Main Results with and without Deprivation

Table B.2 Sub-sample Main Results without Deprivation

	Jewish	**Christians**	**Muslims**
loglincome	0.001	-0.271**	-0.013
	(0.019)	(0.110)	(0.046)
hebrew	0.317***	0.423	0.048
	(0.087)	(0.329)	(0.355)
arabic	-0.168	0.578*	-0.311
	(0.642)	(0.346)	(0.383)
english	-0.815***	-0.601*	-2.057***
	(0.094)	(0.337)	(0.739)
female	-0.046	0.118	0.206**
	(0.036)	(0.162)	(0.097)
academicdegree	0.342***	0.303	0.474***
	(0.045)	(0.195)	(0.127)
married	0.059	-0.024	-0.291***
	(0.047)	(0.197)	(0.113)
age2129	1.920***	2.307***	1.798***
	(0.081)	(0.329)	(0.165)
age3039	1.444***	1.513***	1.225***
	(0.059)	(0.276)	(0.151)
age4049	1.006***	0.954***	0.948***
	(0.056)	(0.263)	(0.152)
age5059	0.428***	0.347	0.389***
	(0.052)	(0.243)	(0.155)
meanloglincome	0.185***	0.771***	0.278***
	(0.030)	(0.165)	(0.070)
cut1	-0.841	2.871	0.316
	(0.267)	(1.317)	(0.691)

Continued on next page

	Jewish	**Christians**	**Muslims**
cut2	0.887	5.112	1.841
	(0.264)	(1.339)	(0.691)
cut3	2.865	7.050	3.239
	(0.266)	(1.368)	(0.693)
σ_u^2	1.295	1.320	0.914
	(0.064)	(0.305)	(0.117)
Observations	14,409	629	2,293
Number of ID	4,592	277	828

Standard errors in parentheses
**** p<0.01, ** p<0.05, * p<0.1*

Table B.3 Sub-sample Main Results with Deprivation

	Jewish	**Christians**	**Muslims**
loglincome	-0.001	-0.302***	-0.016
	(0.020)	(0.114)	(0.047)
dep	-0.095***	0.157	-0.108
	(0.032)	(0.201)	(0.102)
hebrew	0.300***	0.406	0.035
	(0.088)	(0.329)	(0.357)
arabic	-0.167	0.545	-0.288
	(0.645)	(0.345)	(0.385)
english	-0.831***	-0.609*	-2.063***
	(0.095)	(0.338)	(0.746)
female	-0.040	0.101	0.209**
	(0.037)	(0.164)	(0.100)
academicdegree	0.320***	0.309	0.369***
	(0.047)	(0.195)	(0.132)
married	0.068	-0.046	-0.305***
	(0.047)	(0.195)	(0.117)
age2129	1.963***	2.236***	1.779***
	(0.083)	(0.328)	(0.172)

Continued on next page

	Jewish	**Christians**	**Muslims**
age3039	1.472***	1.486***	1.166***
	(0.060)	(0.274)	(0.158)
age4049	1.019***	0.907***	0.891***
	(0.056)	(0.261)	(0.159)
age5059	0.441***	0.354	0.317*
	(0.052)	(0.244)	(0.163)
meanloglincome	0.088**	0.792***	0.133
	(0.043)	(0.220)	(0.088)
meandep	-0.080	-0.154	-0.510**
	(0.060)	(0.398)	(0.206)
cut1	-2.030	2.751	-1.796
	(0.449)	(2.287)	(0.951)
cut2	-0.265	4.956	-0.229
	(0.447)	(2.297)	(0.950)
cut3	1.735	6.871	1.188
	(0.447)	(2.310)	(0.949)
σ_u^2	1.318	1.246	0.943
	(0.066)	(0.297)	(0.122)
Observations	14,209	618	2,234
Number of ID	4,566	275	812

Standard errors in parentheses

**** p<0.01, ** p<0.05, * p<0.1*

Appendix C

Appendix to Chapter 3

C.1 Common Support Check

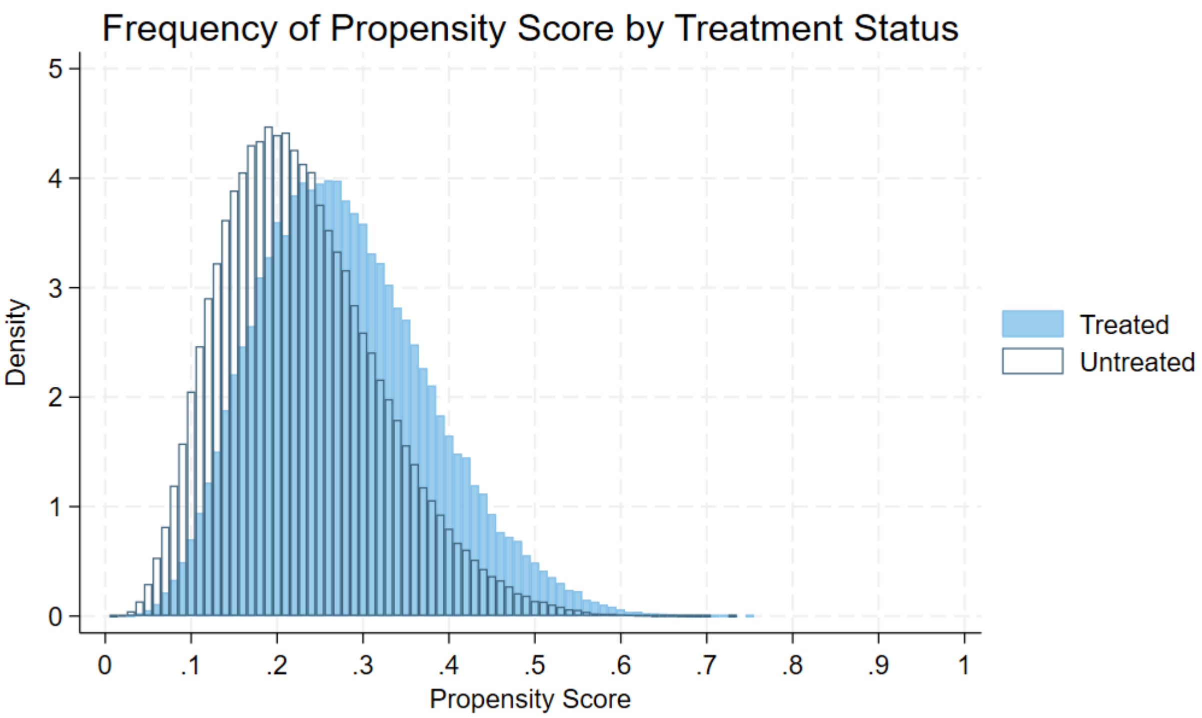

Fig. C.1 Distribution of Propensity Scores by Treatment Status

Figure C.1 displays the distribution of estimated propensity scores for both treated and untreated individuals. The substantial overlap between the two groups suggests that the common support condition is satisfied, allowing for a valid comparison of outcomes across treatment groups. Observations with extreme predicted probabilities (near 0 or 1) are not the case in this analysis. A total of 38 individuals fall outside the common support and are excluded from the analysis.

C.2 Relevence Assumption- GRS

To validate the relevance assumption of the genetic instrument, we regress both BMI-based and waist-to-height ratio measures of obesity on the standardised GRS. The results in *Table C.1* confirm a statistically significant positive association in both cases.

Table C.1 Regression of Obesity Measures on Standardised Genetic Risk Score (GRS)

	Obesity (BMI)	**Obesity (Waist-to-Height Ratio)**
Standardised GRS	0.098***	0.075***
	(0.0008)	(0.0008)
Constant	0.261***	0.214***
	(0.0008)	(0.0008)
Observations	284,270	284,270
R-squared	0.0499	0.0341

C.3 Exclusion Restriction / Validity- GRS

Table C.2 Association Between Covariates and Genetic Risk Score

	Dependent variable: IV GRS	
	Univariate Analysis	Multivariate Analysis
Age	-0.001	-0.001***
	(0.000)	(0.000)
Townsend index	0.008***	0.007***
	(0.001)	(0.001)
Degree	-0.084***	-0.081***
	(0.004)	(0.004)
Female	-0.012**	-0.014***
	(0.004)	(0.004)
Observations	284,308	284,270
R^2	0.001	0.002

Notes: Column (1) reports coefficients from separate regressions of the GRS on each covariate independently. Column (2) reports coefficients from a joint regression including all covariates. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table C.3 IV Estimates of Obesity on Employment: Benchmark vs No-Control Model

	Benchmark (Full)	No-Control	T-statistic	Observations
Obesity	-0.04199***	-0.05404***	1.03	284,270

Notes: Column (1) reports the IV (2SLS) estimate of the effect of obesity on employment controlling for covariates (age, townsend index, degree, female). Column (2) reports the IV estimate without any covariates. Column (3) shows the t-statistic for the null hypothesis that the coefficient in the benchmark model is equal to the coefficient in the no-control model; We cannot reject the hypothesis. Column (4) shows the sample size. '***' indicates significance at the 1% level.

C.4 Robustness Check, Obesity Measured by Waist-to-Height Ratio

We apply the same analysis just with obesity measured by the ratio of waist circumference to height, with values above 0.59 are included as obese. *Table C.4* presents the estimated effects of obesity on the probability of being in paid employment using both LPM and MR estimations. The LPM results indicates a statistically significant negative association between obesity and employment, with an estimated effect of 3.8% reduction points. However, when including GRS for high values of BMI as an instrumental variable for obesity to account for potential endogeneity, the magnitude of the effect increases to 5.5% reduction points, suggesting that the bias in the LPM estimate is likely due to reverse causality or unobserved confounders. These findings follow the results of the main model, when obesity is measured by BMI with the MR approach obtaining a greater effect of obesity on the probability of being in a paid employment.

Table C.4 Comparison of Obesity Effect on Employment Across LPM and MR (2SLS) Models

Variable	**LPM**	**MR (2SLS)**
Obesity Ratio	-0.038***	-0.055***
	(0.0019)	(0.0107)
Age	-0.021***	-0.021***
	(0.0001)	(0.0001)
Townsend	-0.011***	-0.011***
	(0.0003)	(0.0003)
Degree	0.040***	0.039***
	(0.0016)	(0.0018)
Female	-0.003	-0.004**
	(0.0016)	(0.0017)
Constant	1.848***	1.847***
	(0.0066)	(0.0066)
Observations	284,270	284,270
R-squared	0.1153	0.1151
Weak identification test (Cragg-Donald F | First-stage)	—	9720

Note: Dependent variable is being in paid employment. Standard errors in parentheses. *** $p<0.01$, ** $p<0.05$.

The MR model satisfies the relevance condition although the GRS for obesity measured by BMI, as evidenced by a first-stage F-statistic exceeding the conventional threshold for weak instrument detection. The coefficients for control variables such as age, socioeconomic deprivation (Townsend index), and educational attainment remain stable across models, supporting the robustness of the specification.

The updated MTE estimates (*Table C.5*) suggest that the ATE of obesity on the likelihood of being in paid employment is more negative than previously estimated. The new ATE is

–0.053, compared to –0.039 in main analysis. This indicates a stronger negative causal effect of obesity on employment outcomes when observing obesity with waist to height ratio.

Table C.5 Marginal Treatment Effect Results

Variable	**Treated**	**Untreated**
Obesity ATE	-0.053***	-
	(0.013)	-
Age	-0.023***	-0.021***
	(0.0003)	(0.0001)
Townsend	-0.020***	-0.008***
	(0.0006)	(0.0004)
Female	-0.012***	-0.000
	(0.004)	(0.002)
Degree	0.060***	0.035***
	(0.004)	(0.002)
k	-0.016*	-0.007
	(0.009)	(0.008)
Constant	1.866***	1.837***
	(0.021)	(0.006)
Mills: $\rho_1 - \rho_0$	-0.009	-
	(0.012)	-
Number of observations: 284,270		
Treatment model: Probit		

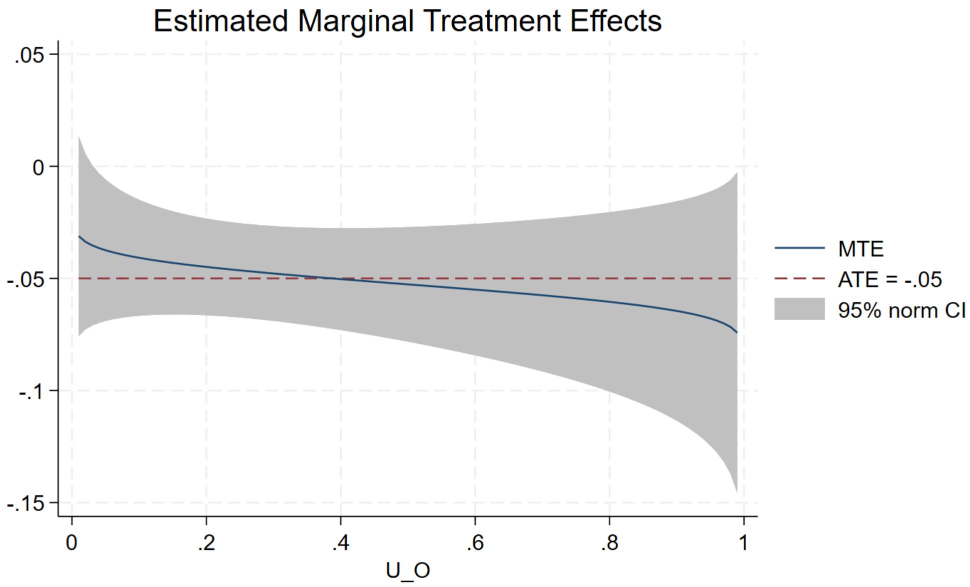

Note: $U_{O_i} = U_{Obesity}$ and represents unobserved factors influencing obesity selection.

When comparing the results of PeT estimates (*Table C.6*), where obesity is measured using the waist-to-height ratio, to the earlier estimates based on BMI, we observe a similar negative pattern in the effect of obesity on paid employment. Here, the marginal PeT effect is also smaller compared to the MR estimates, consistent with the findings from the main model. The ATE remains consistent at –0.028, matching the previous BMI-based result. However, the TT is somewhat more negative at –0.053 (compared to –0.043) following the increase in the magnitude of MR, indicating that individuals classified as obese by waist-to-height ratio may face a greater employment penalty.

Table C.6 Effect of Obesity on Paid Employment (Waist-to-Height Ratio)

	Full Sample	**By Obesity Status**	
	ATE	TT	TUT
Obesity	-0.028	-0.053	-0.022
SD	0.136	0.149	0.132
Observations	284,270	68,478	215,792

Notes: The table reports the mean of *pet_employed* across the full sample and by obesity status, using waist-to-height ratio as the obesity measure.

These subtle variations imply that while BMI and waist-to-height ratio are closely correlated, the latter may capture a slightly different subset of individuals whose employment outcomes are more strongly affected by obesity.

Taken together, the results using waist-to-height ratio as an alternative measure of obesity closely mirror those obtained when using BMI. Across LPM, MR, and the more flexible PeT and MTE frameworks, the estimated effects remain negative and statistically meaningful, confirming a consistent employment penalty associated with obesity regardless of the specific anthropometric definition applied. In both cases, MR estimates yield stronger effects than LPM, suggesting that endogeneity biases downward the naïve estimates.

These parallel findings highlight the robustness of the estimated causal relationship between obesity and employment outcomes. Moreover, they support the use of waist-to-height ratio as a valid alternative or complementary obesity indicator, particularly in sensitivity analyses. While subtle differences, exist especially in the treatment-specific effects, both measures consistently point to a detrimental effect of obesity on labour market participation, reinforcing the broader conclusion that excess adiposity significantly undermines employment prospects.