Slow and Steady: Drawdown Behaviours in Phased Withdrawal Retirement Income Products

Working Paper

UNSW Business School

Igor Balnozan

BActl(Hons)

Supervised by:

Associate Professor Anthony Asher

Scientia Professor Robert Kohn

Professor Denzil Fiebig

Version: January 2018

ABSTRACT

This project analyses twelve years of Australian longitudinal data on drawdowns from phased withdrawal accounts, investigating the behaviours of 44,000 retirees. The dataset used for this analysis is at an aggregate level, based on the combined data obtained from several superannuation funds operating in the industry. First, panel regression models relate drawdown rates to member characteristics. These models indicate the direction, magnitude and statistical significance of the effects of these characteristics on several dependent variables of interest. Second, a cluster analysis allocates members into distinct behavioural groups, characterised by their observed drawdowns over time. Finally, a categorical regression model determines the statistical relationships between member characteristics and the likelihood of belonging to the identified behavioural groups. Although regression models provide some insights into how members draw down their accounts, this project ultimately finds that a small number of simple drawdown strategies explain the vast majority of behaviours within these accounts. Dominant amongst these are two popular rules: adhering to the legislated minimum drawdown rates, and drawing a level dollar amount over time. Many members also make periodic adhoc drawdowns, justifying the need for some flexibility in retirement incomes. To date, the literature has focused on theoretically optimal behaviours derived from lifecycle models. However, a lack of panel data has prevented the empirical observation of these results, as well as a study into the factors which differentiate pensioners into distinct behavioural groups. Consequently, this research bridges the gap between the theoretical results and empirical behaviours. As Australia's legislative environment continues to shift in favour of more flexible arrangements for managing longevity risk in retirement, the findings from this project have important implications for policymakers, financial advisors, and retirement income product designers.

ACKNOWLEDGEMENTS

I would like to acknowledge several individuals whose contributions to this paper have been incredibly valuable.

Foremost, my supervisors Anthony Asher, Robert Kohn and Denzil Fiebig. Thank you for your limitless counsel and guiding support.

Second, I would like to thank Strategic Insight for their tremendous effort in gathering and cleaning a huge quantity of data from several Australian superannuation funds. I'm also indebted to Peter Barnes, who assisted greatly with the preparation of data and preliminary stages of the analysis.

Third, the academics and fellow research students living within the broader School of Risk and Actuarial Studies and at the ARC Centre of Excellence in Population Ageing Research (CEPAR). Amongst these, the contributions from two academics were particularly helpful: Professor Hazel Bateman, who provided extensive feedback on early drafts of several chapters, and Rafal Chomik, as discussant for my presentation at the UNSW Business School Student Research Fair. Both provided a wealth of constructive feedback, particularly in addressing the economic interpretations of results and ensuring adequate coverage of the related academic literature.

Furthermore, two members of the School of Economics were of outstanding assistance in consultations regarding the Econometric components of the methodology. Thank you to postdoctoral researchers Matias Quiroz and David Gunawan.

In addition, I would like to thank the Australian Prudential Regulatory Authority (APRA) and the Reserve Bank of Australia (RBA) for the *Brian Gray Scholarship*. APRA and the RBA established this scholarship program in memory of Brian Gray, who spent more than 20 years shaping public policy in Australia and supporting research relevant to financial institutions. I would also like to acknowledge the financial support provided by the UNSW Business School and CEPAR for their Honours Scholarships.

CONTENTS

Contents **i**

1 Introduction **1**
1.1 Background . . . 1
1.2 Research Motivation . . . 3
1.3 Research Aim . . . 5
1.4 Research Questions . . . 5
1.5 Research Hypotheses . . . 5
1.6 Outline . . . 6

2 Literature Review **7**
2.1 Retirement Incomes and the Decumulation of Wealth . . . 8
2.2 Drawdown Behaviours in Phased Withdrawal Accounts . . . 12
2.3 Advanced Annuity Products . . . 14
2.4 Summary of Literature Review . . . 16

3 Methodology **18**
3.1 Definitions . . . 18
3.2 Data Preparation . . . 19
3.3 Component 1: Panel Regression Models . . . 21
3.3.1 Linear Models . . . 22
3.3.2 'Nonlinear' Models . . . 24
3.3.3 Model Validation . . . 27
3.3.4 Regressors and Regressands . . . 27
3.4 Component 2: Cluster Analysis and Identification of Behavioural Groups . . . 28
3.4.1 Manual Grouping . . . 29
3.4.2 Machine-Assisted Grouping . . . 29
3.5 Component 3: Categorical Regression for Behavioural Group Allocation . . . 31
3.6 Summary of Methodology . . . 33

4 Results **35**
4.1 Component 1: Panel Regression Models 35
4.1.1 Preliminary Analysis 35
4.1.2 Binary Choice Models 39
4.1.3 Continuous Dependent Variable Models 46
4.1.4 Summary of Panel Regression Model Results 60
4.2 Component 2: Cluster Analysis 62
4.2.1 Panel Visualisations 62
4.2.2 Manual Grouping 65
4.2.3 Machine-Assisted Grouping 65
4.2.4 Final Cluster Allocation 74
4.3 Component 3: Categorical Regression Modelling 74
4.4 Other Results and Illustrations 78
4.4.1 Comparing Regular and Adhoc Drawdown Utilisation Rates 79

5 Discussion **84**
5.1 Limitations 84
5.2 Academic Contributions 86
5.2.1 Panel Modelling Contributions 86
5.2.2 Behavioural Contributions 87
5.3 Social Implications 89
5.3.1 Policy 89
5.3.2 Retirement Income Product Design 90
5.3.3 Financial Advice 91
5.4 Future Work 92
5.4.1 Extensions on Available Data 92
5.4.2 Remaining Gaps 92

6 Conclusion **94**
6.1 Addressing Research Aim, Questions and Hypotheses 94
6.2 Summary 96

A Glossary **99**

Bibliography **100**

CHAPTER 1

INTRODUCTION

This chapter motivates a research project that statistically analyses drawdown behaviours observed within account-based pensions—a phased withdrawal retirement income product.

1.1 Background

Retirement systems globally have shifted from Defined-Benefit (DB) to Defined-Contribution (DC) schemes (Broadbent et al., 2006). This move has transferred, from larger institutions to individuals, the management of longevity and investment risk inherent in producing retirement incomes. As a consequence, the modern retiree faces greater responsibility in managing risks and deriving income from their accumulated savings in retirement.

Historically offered by large corporations and the public sector, traditional DB arrangements entitled individuals to an income stream in retirement, usually by reference to a pre-determined formula (Ibid, p3). In contrast, DC funds operate by accepting contributions, generally from individuals or their employers, and subsequently investing strategically to optimise the members' risk-adjusted benefits (Ibid, p7). At retirement, individuals access their accumulated wealth stored in the fund.

Broadly, the income provided by retirement systems is based on three pillars (Bateman et al., 2016). Foremost, governments provide age-contingent—and often means-tested—welfare payments to alleviate poverty in old age. Second, governments may compel or assist individuals to save for their retirement, replacing or reducing the welfare payments required to maintain standards of living in retirement. Finally, voluntary savings behaviour by individuals accumulates assets which can provide additional income in retirement. Often, governments will provide concessional taxation arrangements to incentivise voluntary saving.

Although terminology may differ across countries, the findings of this project are applicable to all countries where DC savings form a significant component of the retirement system. Thus, where ‘superannuation’ or ‘super’ are used in this paper, the terms ‘retirement savings’ or ‘pension system’ may be substituted. In Australia, the first-pillar social security payments are known as the ‘Age Pension’ and are means-tested using both income and asset tests. Approximately 75% of retirees in Australia receive the part or full Age Pension (Bateman and Piggott, 2011).

This project focuses specifically on how retirees utilise second-pillar assets. This encompasses all accumulated capital residing within the tax-favourable superannuation environment. Although mandating second-pillar savings is still relatively uncommon internationally, Australia introduced this policy with the *Superannuation Guarantee (Administration) Act* 1992. Currently, 9.5% of earnings must be contributed to a nominated DC fund. For most individuals, these savings are inaccessible until retirement.

As the superannuation system continues to mature and individuals retire with higher levels of second-pillar assets, Australia represents an ideal case study in the decumulation of superannuation assets—free from selection effects that may exist in DC schemes with no compulsory retirement savings.

Ultimately, the findings from this project must be considered within the wider three-pillar context. For this purpose, Chapter 5 discusses the implications of the results and how they may interact with the other pillars.

Individuals at the threshold of retirement must decide how to access their accumulated assets. Although rules differ between countries, the choice generally lies in: deciding to convert assets into an income stream; retaining control over assets and generating income through the gradual decumulation of available capital; or some combination of the two.

Myriad income products have been suggested to assist retirees in this allocation, both in the academic literature and in practice. Broadly, the purchase of a life annuity can guarantee an income stream for life, while ‘phased withdrawal’ retirement income products assist in drawing down capital throughout retirement. Under this latter arrangement, which are the focus of this paper, an individual remains invested in a combination of risky and safer assets throughout retirement. To generate income, a retiree draws down their account balance over time, possibly subject to annual minimum or maximum rates.

The literature review in Chapter 2 will present a more detailed study of the options available to individuals at retirement. Generally, the allocation of assets between income streams and capital retention involves a trade-off between guaranteed lifetime income and flexible access to accumulated capital.

In Australia, the *Superannuation Industry (Supervision) Regulations* 1994 (SISR) enforces the rules applying to phased withdrawal products. Accounts opened on or after 20 September 2007 are referred to as ‘account-based pensions’, while ‘allocated pension’ is used to describe similar products existing prior to 2007. Throughout, we refer to these older products as ‘legacy accounts’, and accounts opened more recently as ‘modern accounts’. Annual drawdowns from both account-based and allocated pensions must adhere to minimum rates, as specified in the regulations. Additionally, before 1 July 2007, legacy accounts were subject to maximum drawdown requirements.

1.2 Research Motivation

Historically, the retirement income stream products available in Australia have been restricted to traditional guaranteed term and lifetime annuities (Asher, 2015). Restrictive regulations have barred more advanced variants, such as deferred or variable annuities, from entering the Australian market. Moreover, the market for traditional annuity products was virtually non-existent in the early 2000s (Bateman and Piggott, 2011). Chapter 2 will describe many possible explanations for this lack of annuitisation from both the supply and demand side—as well as some evidence for slight growth in annuity demand during recent years.

Recently, the Australian government has removed regulatory obstacles to the development of more advanced retirement income products, such as variable and deferred annuities (Australian Government The Treasury, 2016a). Largely, this has been in response to appeals presented in the academic literature and policy research such as the Henry Review (Commonwealth of Australia, 2010). Having removed the regulatory barriers, the government intends to begin promoting hybrid products, referred to as Comprehensive Income Products for Retirement (CIPRs) (Australian Government The Treasury, 2016a). CIPRs will aim to promote longevity insurance by offering income streams in combination with the liquidity and investment freedom available in phased withdrawal products.

The development of suitable CIPRs requires inquiry into retiree drawdown behaviours, in order to understand their preferences for income, risk-management, and flexibility, which have been identified by the government as competing objectives. For instance, the design of appropriate benefit structures for retirement income products should consider whether desired income in retirement is level, increasing, or decreasing. Moreover, three-quarters of Australians regularly receive Age Pension payments, which form a longevity- and inflation-protecting income stream. Correspondingly, we can look to drawdown behaviours within account-based pensions to identify whether individuals use their liquid second-pillar assets to create their own income streams above the minimum drawdown rates. This would identify individuals who desire income streams above and beyond those already guaranteed by the government—or those who are too wealthy to obtain the Age Pension. Furthermore, examining the extent to which retirees use their account-based pensions to make adhoc drawdowns in retirement can help determine appropriate recommended levels of precautionary savings.

Academic research in the field of behavioural economics has underscored the impact of default options in decision-making (see e.g. Kahneman, 2003, p1459). Furthermore, recent work by Bateman et al. (2017) highlights that these findings are indeed applicable to individuals making financial decisions at retirement. Briefly, the findings imply that options which are given default values in product design, such as asset allocations and benefit amounts, will be gravitated towards in decision-making. In this respect, the design of default options for CIPRs requires considerable empirical research, as these defaults may determine what the majority of individuals choose at retirement. For example, within a CIPR, there may be multiple parameters to set at retirement, including a base level of income guaranteed by annuitisation, and a minimum amount of precautionary savings insured at any given time. Without defaults on these parameters, the process of committing to a CIPR may be daunting and require substantial financial advice, whereas with appropriate defaults, the majority of retirees may be automatically guided towards making prudent decisions. Moreover, the decision to select a CIPR at retirement could itself become the default option within a superannuation fund, and hence there is a considerable burden on the government and super funds to design these appropriately to meet the needs of retirees—specifically, those needs identified in the literature and in empirical analysis.

In the literature on phased withdrawal income products, Bateman and Thorp (2008) evaluate several retirement drawdown strategies that might be utilised by rational individuals, comparing them to simulated optimal behaviours. However, these results did not include the impact of the Age Pension on optimal drawdown behaviours. Empirical panel analysis of observed drawdown behaviours has the potential to extend the literature on optimal decumulation, by finding deviations from theoretical results and identifying the ways in which precautionary savings are used.

In the related literature, Asher et al. (2017) study the decumulation of total assets in retirement using panel data on a sample of retirees receiving the Age Pension. Furthermore, work by Hulley et al. (2013) and Spicer et al. (2016) similarly investigates the movement of retirement assets using the longitudinal HILDA dataset.

A gap that has remained in this literature is understanding the drawdown behaviours within phased withdrawal accounts specifically, rather than the decumulation of total assets. Due to the effect of asset allocations and investment returns on balances in account-based pensions, the true second-pillar asset drawdown behaviours have not been visible to previous researchers. As a result, our understanding of retirement decumulation behaviour has been incomplete.

Despite these motivations, existing empirical analysis on drawdown behaviours has been inadequate in meeting the above literature and policy needs, partly due to a lack of relevant longitudinal data. Poterba et al. (2013) analyse data on withdrawals from personal retirement accounts in the United States, but a lack of panel data and an inability to distinguish between regular and adhoc drawdowns limit the applicability of their results. However, large longitudinal datasets for both APRA-regulated and self-managed super funds have recently been collected in Australia, remedying the above issues. Currently, only descriptive analyses have been conducted on these datasets (see both Sneddon et al., 2016; Plan For Life, 2016), while this paper applies statistical methods in analysing the data.

In summary, the development of retirement savings systems and concomitant financial products is in a transitional phase. Globally, retirement systems are shifting from simpler arrangements, underwritten by larger corporations, to more advanced solutions, requiring retirees to take on greater responsibility for risk-management in retirement. Crucially, to design better retirement income products and provide appropriate advice, it is important to understand the empirical behaviours observed by individuals choosing to generate income in retirement through phased withdrawals from an investment-linked account. Concurrently, the literature on the decumulation of second-pillar assets is lacking a statistical analysis of the empirical drawdown data. Achieving this would enrich the empirical literature on retirement decumulation, and further bridge the gap between theory and reality. These needs, arising from contextual factors and the related literature, serve to motivate this project.

1.3 Research Aim

Consequently, this project aims to:

Identify and explain drawdown behaviours in phased withdrawal products

The impact of fulfilling this aim is two-fold. First, it will progress the academic literature on drawdown behaviours within phased withdrawal accounts, which until now has relied primarily on theoretical studies into optimal behaviours, and lacks feedback from empirical studies. Second, it will provide timely insights into appropriate policy decisions, retirement income product design, and financial advice, during a transitional period for Australia's retirement system.

1.4 Research Questions

In fulfilling the above aim, the research must address the following three questions:

1. What drawdown behaviours are observed in account-based pensions?
2. Are statistical models effective at predicting drawdown rates and behaviours?
3. Which income products and policy design recommendations would suit the identified groups of retirees?

1.5 Research Hypotheses

Existing research suggests the following hypotheses to test throughout this project.

Annual Drawdown Rates

1. Older individuals draw down less in excess of the minimum rates, compared to younger retirees
2. Individuals with larger account balances draw less in excess of the minimum rates, compared to retirees with smaller account balances

3. Females draw more slowly through their account balances than males, after controlling for factors such as account balances
4. In financial years following the GFC, drawdowns in excess of the minimum rates decreased
5. In financial years following the GFC, the temporarily lower (concessional) minimum drawdown rates encouraged many retirees who had been drawing at the previous minimum rates to reduce their drawdowns to the concessional levels

Behavioural Groups in the Drawdown Series

1. A substantial portion of retirees will draw consistently at minimum rates
2. A group will attempt to draw at a constant rate, for example 7% per year
3. Some will draw a constant nominal—not rising with inflation—dollar amount throughout retirement
4. A group will draw a constant real—rising with inflation—dollar amount
5. Some retirees will spend more than the minimum rates initially, but over time reduce drawdowns

1.6 Outline

The remainder of this paper has the following structure. Chapter 2 reviews the related literature, exploring how existing work interacts with this project and explaining the gap this research aims to address. Subsequently, Chapter 3 describes the methods used to analyse the available data, while Chapter 4 presents the research findings. Chapter 5 discusses these results in depth, highlighting the academic contributions and social implications of the findings. Finally, Chapter 6 links the results back to the research aim, questions and hypotheses, and provides a summary of the paper.

CHAPTER 2

LITERATURE REVIEW

This chapter explores key papers from the literature on retirement decumulation theory and practice, highlighting the gap this project aims to address. We begin by investigating the decumulation of assets in retirement more broadly, before honing in on drawdown behaviours within phased withdrawal accounts specifically. It is with this focus on second-pillar asset decumulation that we progress the academic literature. Finally, as the contextual motivating factors for this research include the design of more appropriate income streams in retirement, we also review the literature specific to the design of annuity products. While traditional guaranteed term and lifetime annuities are regularly available, the uptake of more advanced variants has been staggered around the world. Australia in particular has, until now, offered a restrictive environment for the development of these products, however recent legislative changes have removed these barriers.

In digesting such a broad literature, it will be helpful to borrow and extend the terminology used by MacDonald et al. (2013). Research into decumulation in retirement answers at least one of four questions: 'How Should?'; 'How Could?'; 'How Can?'; and 'How Do?'.

The phrase 'how should' will be used in reference to studies deriving theoretical, optimal behaviours that rational retirees are conjectured to exhibit. Often, a utility function underlies the derivation of these behaviours. As will be shown, these have expanded from the simpler assumptions used by Yaari (1965), which lent themselves to closed-form solutions, to more sophisticated models (e.g. in Iskhakov et al., 2015), which require simulation analysis to arrive at conclusions.

'How could' refers to the design of income products that retirees potentially could utilise in their financial decision-making. Generally, these are products that have been proposed by the literature, but are not yet offered by superannuation funds or other financial institutions.

In contrast, 'how can' describes products that are already available in the retirement incomes market. As examination of recent legislative changes in several countries will show, the 'how can' component contains a strong interaction with public policy.

Lastly, the literature investigates 'how do' retirees draw down their wealth, through statistical analysis of the empirical data on financial decisions observed in retirement. Importantly, MacDonald et al. (2013) note that the 'how do?' question has had the least attention in published papers, despite being necessary to feed back into the 'how should' and 'how could' literature. One explanation for this lack of 'how do' papers is the possible difficulty in collecting the requisite data.

2.1 Retirement Incomes and the Decumulation of Wealth

Retiree annuitisation—converting a lump sum into an income stream—has been encouraged to varying degrees by governments around the world. Some countries mandate partial or full annuitisation of DC savings at retirement, while retirees in other countries have complete autonomy over their finances—although some countries, like Australia, offer incentives for annuitisation (Mercer, 2016, p56).

These arrangements exist in a state of flux, indicating that countries are not unanimous in their response to contemporary concerns. For example, Singapore has introduced compulsory annuitisation within the last decade (Fong et al., 2010), while the United Kingdom, previously requiring annuitisation by age 75 (Emms, 2010, p176), has recently removed this regulation (Mercer, 2016, p47).

In Australia, legislation historically restricted the design of income products other than traditional guaranteed annuities and account-based pensions (Stringer, 2011; Clare, 2013). Until recently, SISR sections 1.05 and 1.06 required that account balance products must pay at least the legislated minimum drawdown rates, and that income stream products must pay a predetermined amount annually for the life of the holder (Stringer, 2011; Australian Government The Treasury, 2016b, p9).

In response to continued criticism, the Australian government committed to effecting meaningful change in its 2016 Federal Budget. In particular, they proposed the removal of barriers to freedom in the design and implementation of more advanced annuity and pension products (Australian Government The Treasury, 2016b). These legislative changes have since been enacted, and as noted in Chapter 1, the quest for superior income products is supported by the Australian government through its development of CIPRs.

While legislators have engaged with the influence of regulations on the breadth of the annuity market, the academic literature has continued to explore retiree attitudes towards annuitisation. In particular, a commonly quoted conundrum is the low levels of voluntary annuitisation observed in most countries throughout the second half of the 20th century. To contextualise the issue, it should be noted that one of the contributions of Yaari (1965) was to conclude that a rational individual—conforming to several restrictive utility assumptions, including a lack of bequest motive—should convert all wealth at retirement into a guaranteed lifetime annuity. Brown (2009) and Benartzi et al. (2011) cite Modigliani's 1985 Nobel Prize acceptance speech, revealing that the annuity 'puzzle'—a shortfall in annuitisation behaviour, relative to the expectations from the literature—has been known for several decades. This dearth in annuitisation continued to be observed after the turn of the century: in the United States (see e.g. Mitchell et al., 1999; Brown, 2009); in Australia (Bateman and Piggott, 2011); and generally, around the world (see e.g. James and Song, 2001).

A wide range of rational explanations for this departure from the original theory is summarised by Brown et al. (2008). Broadly, these justifications conclude that traditional annuity products may not meet the needs of a rational retiree due to a variety of possible factors, including: bequest motives; the need for liquidity and precautionary savings in retirement; and pre-existing annuitisation provided by public welfare systems.

Research into rational reasons for low annuity demand has shown that these do not entirely explain annuitisation behaviour. For example, Lockwood (2012) compares the results of a simulation study with empirical data to determine the validity of the bequest motive as a determinant of annuitisation decisions. Whilst the simulation results imply that bequest motives should be a significant factor, when analysing the data, individuals with strong bequest motives exhibited very similar rates of annuitisation to those with weak bequest motives. Lockwood notes that this is broadly consistent with the findings of Brown (2001), discovering that individuals self-reporting a higher level of importance placed on leaving a bequest did not annuitise their DC fund balances at significantly lower rates.

Indeed, Brown (2009) suggests that the explanations for low annuitisation rates need not assume individuals are behaving rationally. Some behavioural hypotheses for low annuitisation levels explored are:

- The framing effect—whether annuities are presented as an investment decision, or a consumption guarantee
- Complexity and financial literacy—Lusardi and Mitchell (2007) show evidence that retirees do not have the financial literacy required to deal with the complex financial decisions retirement now presents
- Mental accounting and loss aversion—individuals consider an annuity as wasted capital in the scenario where they die younger than expected
- Misleading heuristics—individuals view insurance as protection from adverse outcomes, but struggle to consider living 'too long' as the adverse scenario in retirement
- The illusion of control—believing that retaining control over one's assets will improve financial security in the future

Brown et al. (2008) study the framing hypothesis by presenting participants with actuarially equivalent choices, differing only in their framing. Their results found that framing annuities as investment decisions reduces their appeal, whereas presenting them as consumption guarantees makes them more attractive.

Considering the literature on behavioural impediments to annuitisation, and despite the low levels observed, Benartzi et al. (2011) ultimately find evidence that individuals value and even prefer annuities—when the underlying conditions are right. In addition to having enough accumulated capital to make an annuity purchase worthwhile, the option must be presented at the appropriate age, and with the right framing.

Extending this literature, Beshears et al. (2014) study survey data on hypothetical annuity purchase decisions. The results suggest that partial annuitisation is preferred to complete or no annuitisation, and that products which can provide additional choice and flexibility are more popular—for example, an annuity which provides a bonus payment during one month each year. Furthermore, they confirm the findings of Brown et al. (2008) regarding the significance of the framing effect: ignoring the implied investment returns generated by annuities increases their appeal.

Thus there have been attempts to justify low annuity demand using rational reasons, as well as by investigation of relevant cognitive biases. At least in Australia, however, the literature has identified evidence of an increase in annuity demand in recent years (Iskhakov et al., 2015, citing Plan For Life, 2014). Indeed, an inspection of recent annual and interim reports issued by Challenger, one of Australia's life insurers, evidences growth in annuity sales (see e.g. Challenger Limited, 2016, 2017a).

The culmination of these theoretical studies and empirical observations has been an expansion of the factors considered in the decumulation phase by theoretical lifecycle models. For example, Iskhakov et al. (2015) complete a comprehensive analysis of optimal annuitisation under a range of scenarios, by running simulations against a more sophisticated stochastic lifecycle model. One key contribution from this paper was a consideration of how access to means-tested social security payments—in Australia, the 'Age Pension'—crowd low-wealth households out of the annuity market completely.

Moreover, research by Bateman et al. (2017) has engaged directly with the hypothesis that cognitive biases influence financial decisions in retirement. As mentioned in Chapter 1, this study highlights the significant role that default options play in guiding the choices individuals make with regard to annuity choice. Specifically, Bateman et al. find that when allocating assets between life annuities and account-based pensions, individuals generally prefer some combination of the two. Crucially, certain types of people prefer to stick with the default allocation presented to them, while others follow simple heuristics (rules of thumb): either a 0–100%, 50–50% or 100–0% split. In this project on drawdown behaviours in phased withdrawal accounts, we similarly investigate the impact of default options and simple heuristics on behaviours—within account-based pensions, specifically.

Goda and Manchester (2013) draw similar conclusions regarding the powerful effect of default options in determining the choice of retirement fund. Where individuals are given a choice between a DB or DC fund to accumulate wealth for retirement, encountering a DC plan as the default option made individuals 60% more likely to 'choose' the DC fund.

Although DC schemes are becoming the standard for the accumulation phase of retirement, these increase the risk-management responsibilities of individuals—and not only during the decumulation phase. Ganegoda and Evans (2017) develop an economic scenario generator to observe the impact of shocks to investment returns within DC accumulation funds, specifically accounting for the possibility of low-frequency, high-impact market shocks. Depending on the timing of retirement relative to these shocks, individuals with otherwise similar accumulation behaviours can experience large differentials in wealth at retirement. To protect against this downside risk, they recommend option-like portfolio insurance strategies within the accumulation funds.

Furthermore, second-pillar asset decumulation behaviour is closely linked to the means-testing applied in determining eligibility for first-pillar safety nets such as Australia's Age Pension. Hulley et al. (2013) study this interaction by first simulating optimal asset decumulation strategies in retirement in the presence of the means-tested Age Pension, and subsequently investigating the empirical experience using the longitudinal HILDA dataset. Simulations suggest that individuals who are close to or within the eligibility criteria for receiving means-tested public pensions should decumulate their assets faster, and place a higher proportion of their wealth into risky assets. In this way, they maximise their entitlements to first-pillar income, while the government underwrites their private asset investment risk. Indeed, empirical data analysis confirms these theoretical results. Moreover, decumulation overall occurs at modest levels for less well-off retirees, while wealthier retirees tend to accumulate in early retirement by adoption of riskier investment strategies. Evidence emerging from the United Kingdom also shows that, far from beginning to decumulate immediately, more than 75% of individuals continue to increase their savings after retirement (Brancati et al., 2015).

In the same stream of literature, Asher et al. (2017) apply regression models to longitudinal data from Australian social security—'Centrelink'—payments to 10,000 Age Pension recipients. Overall, consumption appears to have been conservative, with a majority of pensioners passing on significant bequest sums on death. Further, the data show that consumption declined with age, instead of increasing in line with expectations of rising medical costs. Many pensioners even continued to accumulate in the early stages of retirement, a finding which resonates with that of Hulley et al. (2013) and Brancati et al. (2015) above. Asher et al. conclude that if bequest and precautionary motives are ignored, most pensioners could currently afford to spend more without exhausting their savings during retirement.

Critically, the Centrelink dataset is subject to a selection effect, only sampling from individuals receiving welfare payments. Additionally, the treatment of superannuation assets by the sampled individuals was not visible, which is the focus of the present study. Consequently, it will be insightful to compare the findings of this project, utilising a panel dataset on account-based pensions, with the findings from social security recipients above.

Interestingly, the effect of health and ageing shocks on retirement wealth depends greatly on country-level effects. This is made clear on comparing two similar studies on the evolution of household wealth throughout retirement: in the US by Coile and Milligan (2009); and in Australia by Spicer et al. (2016). In the US, the effect of health-related shocks has a significant impact on retirement wealth, and results in retirees liquidating housing and other assets. In contrast, Australians are impacted much more lightly by shocks to health, attributed to a more generous subsidised public healthcare system. Furthermore, Australians prove very reluctant to release housing wealth. Spicer et al. (2016) note also that a Dutch study by Van Ooijen et al. (2015), in which retirees face a similarly generous healthcare system, mirrors the Australian case, rather than that of the US.

2.2 Drawdown Behaviours in Phased Withdrawal Accounts

Of direct interest to this project is the existing body of theoretical and empirical work investigating the drawdown of second-pillar assets, especially within phased withdrawal products. Specifically, analysis of post-retirement drawdown habits is crucial for informing better financial management, product creation, and government policy design (Plan For Life, 2016).

Studies into how retirees 'should' draw down from phased withdrawal products can be traced back at least to the second "Pensionmetrics" paper from Blake et al. (2003). This paper was strongly tied to the regulatory situation effective at the time in the United Kingdom. Specifically, by age 75, retirees were required to annuitise any remaining balance within their DC fund accounts. Between regular retirement and age 75, individuals were able to use their retirement savings with greater freedom, including the ability to open a phased withdrawal account, similar to Australia's account-based pensions.

Blake et al. (Ibid) compare three options for an individual retiring at age 65: purchase of a level annuity for life; purchase of an investment-linked annuity until age 75, at which time the remaining value was converted to a lifetime annuity; or opening a phased withdrawal account and drawing down to generate income until 75, when the balance would be similarly converted to a lifetime annuity. Instead of searching for optimal drawdown strategies within the phased withdrawal product, however, the paper concludes that broadly equivalent outcomes can be generated within each of the three options considered, by varying the individual's exposure to equity returns in retirement. Additionally, within the latter two options, the age to annuitise is varied in an attempt to find the optimal annuitisation age. However, the observation that risk appetite—the willingness to expose oneself to risky returns—determines one's behaviour is important, and is a factor which we consider in this project.

With an increased interest in phased withdrawal products specifically, Horneff et al. (2008) adopt a utility-based framework with stochastic return rates and retiree lifetimes, to compare three drawdown strategies alongside the level payments implied by a guaranteed lifetime annuity. The strategies include: drawing a fixed proportion of the account balance annually; drawing a proportion equal to $1/T$, where T is defined as the theoretical maximum remaining lifetime; and drawing $1/E[T]$, where $E[T]$ is the new remaining life expectancy at each surviving year. Ultimately, Horneff et al. reposition their findings to seek the optimal age to annuitise, which is of less interest to the current project. An introduction of their aims, however, is instructive before reviewing the superseding work by Bateman and Thorp (2008).

Bateman and Thorp, similarly considering the above strategies within a stochastic lifecycle model, extend the work of Horneff et al. by including, as competing strategies, the newlylegislated set of minimum drawdown rates applying to account-based pensions in Australia. These begin at 4% for retirees under 65, and increase progressively with age in seven stages, reaching a maximum of 14% for individuals 95 and older. As a result, each strategy, including drawing at exactly the Australian minimum drawdown rates, could be directly compared with optimal drawdown behaviours, derived by simulation from the assumptions placed on an individual's utility function. On balance, the authors find that the legislated minimum drawdown rates are relatively close to the optimal behaviours. In many scenarios, however, a fixed-percentage drawdown rule increases simulated utility. Hence the literature on optimal drawdown behaviours in phased withdrawal products progressed substantially in the 2000s.

Despite this progress, a similarly substantial branch of literature into the empirical experience in phased withdrawal accounts has not yet emerged. Some attempts have been made, however these have not been adequate in fulfilling the needs outlined in Chapter 1, arising from the literature and from policymakers. Furthermore, there remains a gap where one would expect research providing the necessary link between the theoretical literature and reality.

Perhaps the attempt which has come the closest is the research by Poterba et al. (2013) into drawdowns from personal retirement accounts in the United States. Poterba et al. ran several statistical models to fit various dependent variables in the observed data. These included not only the drawdowns as both dollar amounts and proportions of account balances, but also binary choice models to estimate the probability that an individual makes any drawdown, in years where this is not compulsory.

Critically, the research by Poterba et al. was limited by two key factors. Firstly, the data available did not observe individuals over the duration of their sample period, and so several cross-sectional or shorter-panel datasets through time were pooled to create a "synthetic" panel (p7). As a result, the methods employed were unable to control for any unobserved heterogeneity in drawdown behaviours of individuals. Furthermore, drawdown behaviours, as they are defined in the context of this project, are observed over time, and not solely at one point in time. Achieving this research goal requires a panel dataset, tracking individuals over longer time periods. As will be detailed in the methodology, the panel dataset utilised by this project is an advancement in this respect. Secondly, the results of Poterba et al. do not differentiate between 'regular' drawdowns, which are nominated to be received over time as an income stream, and 'adhoc' drawdowns, which an individual can commute from their account balance to meet larger or unexpected costs. This desirable feature is another characteristic of the newly-available data. Consequently, it is argued that the literature requires a paper to fill the gap left by Poterba et al.

To the best of our knowledge, since 2013 there has not been a statistical attempt to complete this stream of the literature. Recently, a longitudinal dataset has been made available, but to date, only descriptive analytics have been performed on it, by Plan For Life (2016) and Sneddon et al. (2016). The former considers the data from APRA-regulated funds, while the latter analyses the data on self-managed super funds.

The Plan For Life report on superannuation fund data showed that in approximately 50% of cases, drawdown was done at the minimum level. The report also found that in the year preceding death, drawdown often became rapid and unsustainable, possibly to fund out-of-pocket medical expenses, suggesting a need for more long-term health and longevity insurance solutions. Notably, Plan For Life recognised the need for further work to be carried out on their data. Broadly, Sneddon et al. mirror the findings above, with most retirees in their 60s and 70s drawing close to the minimum amounts each year.

Despite reporting on the aforementioned panel dataset, the above sources lack a rigorous statistical methodology, instead limiting the analysis to descriptive statistics and summary data. Furthermore, there has been no attempt to exploit the panel nature of the data to identify patterns in the drawdowns over time. The benefits of estimating statistical models are two-fold: it is possible to conduct robust inference on the statistical significance and signs of the parameters corresponding to all observed characteristics of the individuals; and models which prove successful at predicting out-of-sample results can be used to estimate the drawdown behaviours of retirees not captured in the panel. Hence our project remedies this gap in the empirical literature.

Consequently, this research contributes to two streams of literature. The first is the theoretical literature on drawdown behaviours in phased withdrawal accounts, which this research extends by exploring how observed drawdown patterns relate to the theoretical results. Second, the findings from this project complement other work in the empirical literature on the decumulation of wealth in retirement, including the Centrelink data analysis by Asher et al. (2017), as well as analysis of HILDA data by Hulley et al. (2013) and Spicer et al. (2016). Where these other studies have been unable to observe the rates at which retirees draw down their second-pillar assets within phased withdrawal accounts, we study this aspect of decumulation specifically. As a result, a richer view of the financial experience of retirees in Australia emerges.

2.3 Advanced Annuity Products

This section of the literature review serves to construct an image of what a developed market for retirement income products might resemble. In particular, one question underpins all the following papers: in theory, how 'could' retirees generate income from their accumulated wealth?

While the design of advanced retirement income products in Australia has been restricted in the past, other countries have successfully been using advanced products to manage the risks and meet the financial requirements of retirees, especially in the US, Asia and Europe (Asher, 2012; Clare, 2013; Institute of Actuaries of Australia, 2014). The Institute of Actuaries of Australia outlines the defining characteristics of several of these proposed solutions, including: Pooled Annuities and Group Self-Annuitisation Products (GSAs); Guaranteed Lifetime Withdrawal Benefit (GLWB) riders on Variable Annuities (VAs); With-Profit Annuities (WPAs); and Deferred Lifetime Annuities (DLAs). Three other noteworthy product designs, not included in the Actuaries Institute review, are: the Life Care Annuities suggested by Wu et al. (2016); the Longevity-Indexed Lifetime Annuities proposed by Denuit et al. (2011); and the Longevity-Indexed Deferred Annuities also from Denuit et al. (2015). The remainder of this section will provide an overview of the characteristics of these proposed products.

Qiao and Sherris (2013) extend the idea of GSA schemes introduced by Piggott and Detzel (2005), providing solutions to some shortcomings of the initial presentation. GSAs allow individuals to pool capital up-front, and use this capital to make regular annuity payments to surviving members, while funds suffice. In their original specifications, GSAs suffered from the limitation that as the pool matured, its size naturally shrank due to the death of self-annuitants. Correspondingly, the reduced pool size increased the variability of the received payments over time. Critically, one of the main motivations for annuitisation—longevity insurance against outliving savings—was undermined by GSAs, as the longest-surviving pool members became increasingly likely to exhaust the funds in the pool in the presence of high longevity experience.

Notably, Qiao and Sherris use simulated pool dynamics to suggest two simple improvements. First, the authors show that increasing the pool size is very effective at reducing the late-life benefit payment volatility. Second, and more significantly, allowing new cohorts to join the pool after commencement of the original scheme has a similar effect in the reduction of payment volatility for the longest-surviving members, and reduces the expected drop-off in benefit payments in the presence of improving longevity.

The contribution of Donnelly (2015) was to provide a detailed comparison of the Group Self-Annuitisation (GSA) scheme, the Pooled Annuity Funds (PAFs) of Stamos (2008), and the Annuity Overlay Fund (AOF) of Donnelly et al. (2014), which achieve similar risk-sharing goals through different mechanisms. In particular, Donnelly highlights conditions under which actuarial fairness is attainable for each style of annuity product, which serves to increase the desirability of the product to consumers.

As an alternative to risk-sharing by the pooling of funds by individuals, payments from an annuity provider can be indexed in reference to relevant characteristics. Existing papers by Denuit et al. (2011) and Richter and Weber (2011) argue that indexing variable annuity payments to longevity trends is one solution in managing longevity risk. Under this arrangement, some or all of the systematic risk component is shared between the insurer and the annuitants. Importantly, the annuitant still retains protection against outliving their assets, but benefits from a lower product cost due to the insurer's reduced capital requirements.

Denuit et al. (2015) also explore the impact of indexing the deferment period on longevity products such as deferred life annuities and reverse mortgages. In effect, this makes the sharing of longevity risk an intra-, rather than inter-, generational cost, with the insurer bearing interest rate and any idiosyncratic risks, and annuitants pooling their systematic longevity risk.

In contrast, annuity benefits can instead by indexed to the investment performance of a reference portfolio, allowing annuitants with higher risk appetites to link their benefit payments to the returns of risky assets.

Milevsky (2013) provides a comprehensive overview of this type of investment-linked variable annuity. Where sold, investment-linked annuities compete directly with phased withdrawal account-based products. While both investment-linked annuities and phased withdrawal accounts allow the satiation of risk appetite—resulting in periods of greater consumption when investment returns are favourable (and vice-versa)—investment-linked annuities forego access to a larger, liquid stock of wealth, in favour of guaranteed longevity insurance.

An individual need not necessarily choose only one of these two desirable features, however. A rider—an optional 'add-on'—increasingly common to investment-linked annuities are known as GLWBs—Guaranteed (Minimum) Lifetime Withdrawal Benefits (Ibid). For an additional cost, these riders insure a minimum level of liquid capital that can be accessed throughout the duration of the contract, creating a product which forms a compromise between investment-linked annuities and phased withdrawals.

Finally, in countries where healthcare and long-term care expenditure is insufficiently subsidised by the government, these costs may be a significant motivator for conservative consumption in retirement (De Nardi et al., 2015; Wu et al., 2016). A product proposed by Wu et al. is the 'Life Care Annuity', which combines the benefits of a traditional guaranteed lifetime annuity with insurance against late-life healthcare expenditure. The results indicate that this specification would increase the attractiveness of annuitisation, although the impact is contingent on the adequacy of a country's public healthcare system.

Consequently, following recent legislative changes, the Australian superannuation system is well placed to benefit from the design and implementation of more advanced retirement incomes solutions. Globally, nations are at different stages in the development of decumulation options and retirement income product markets. As these markets continue to mature, the literature surveyed suggests myriad products tailored to meet the heterogeneous needs of individuals in retirement.

2.4 Summary of Literature Review

The relevant literature on retirement savings and spending answers four questions: how retirees should, could, can, and do, draw down on their accumulated wealth. Moreover, this chapter has made clear that none of these streams exist in isolation. Instead, there is a complex interplay between all four questions.

Papers in the literature on optimal behaviours—'how should'—can be motivated by empirical observation—'how do'—or by government policy and the resulting development of financial markets for relevant insurance products—'how can'. The findings from the optimality literature, however, require the collection of richer data to test new hypotheses and identify the deviations from results derived by simulation against utility frameworks.

Legislation may take time to adapt to the rapid pace presented by the literature, but this conservative position may protect individuals from false positives or misconstrued results. Notably, different countries contend with a diverse range of contextual factors, and it is clear that responses to the challenges of population ageing are contentious and equally varied.

Ultimately, this chapter shows that a key gap in the literature remains to be filled. To date, theoretical work has studied optimal behaviours in phased withdrawal accounts in isolation from the impact of the Age Pension. Until now, empirical studies on retiree drawdowns from phased withdrawal products have been unable to provide adequate insights into the true behaviours within these accounts. Understanding these behaviours is critical as policymakers and financial product designers continue developing the menu of financial options available in retirement. In particular, they require a better understanding of how individuals prefer to draw down their second-pillar assets, which the existing decumulation literature has been unable to provide. An area of interest is the extent to which retirees need the flexibility of holding reserves of liquid capital while still deriving a stream of income, as phased withdrawal products allow. Moreover, CIPRs may contain default options—for example, regarding the allocation of superannuation assets to income streams and precautionary savings. Due to the power of defaults in gravitating individuals towards predetermined options, it is critical that these defaults be informed by empirical data.

CHAPTER 3

METHODOLOGY

This chapter presents three methodological components—panel regression models, cluster analysis and a categorical regression—that make effective use of the large panel dataset available. The first component focuses on how retiree characteristics influence drawdown rates at individual points in time; the second component segments panel members into distinct behavioural groups; and the third component investigates the characteristics that are significant in determining which group an individual belongs to.

3.1 Definitions

In this study, 'drawdown' refers to the withdrawal of account value—measured over a complete financial year. We classify drawdowns along three dimensions:

1. Amount or Rate
2. Nominal or Excess
3. Regular or Adhoc

Drawdown 'amounts' are the dollar figures withdrawn from the phased withdrawal account. The corresponding 'rates' are calculated by dividing the amount drawn down by the account balance at the beginning of the corresponding financial year.

$$\text{Drawdown `rate'} = \frac{\text{Annual drawdown amount}}{\text{Account balance at financial year start}} \tag{3.1}$$

This convention for calculating annual drawdown amounts and rates aligns with the method used to determine the minimum annual drawdown requirements, as specified in the SISR.

These nominal amounts and rates must satisfy the legislated minima. We define 'excess' drawdown as the difference between the nominal drawdown and the corresponding minimum required. Note that since 1 July 2007, there has been no upper limit on the drawdown rate.

Prior to 1 July 2007, the minimum drawdown rate changed for each year of age up to age 100. SISR schedules 1A and 1AAB contain these tables. In 2007, the government simplified the minimum drawdown rates, leaving the more parsimonious rules contained in Table 3.1.

As made evident in Chapter 4, during financial years 2008 and 2009, adverse economic conditions significantly eroded account balances. To ameliorate the impact on retiree savings held in account-based pensions, the government introduced concessional minimum drawdown arrangements for several years following. In financial years ended 30 June 2009–11 inclusive, the concessional rates were 50% of the usual rates, while in financial years ended 30 June 2012 and 2013, the concessional rates were 75% of the usual rates. For example, a retiree aged 65 on 1 July 2012 faced a minimum drawdown rate of 3.75% for financial year 2013. The following year, their minimum drawdown rate was 5%.

Finally, retirees can nominate, in advance, the amounts and frequencies of the payments to be drawn from their account-based pension. We refer to this prospective drawdown allocation as the 'regular' drawdowns. One of the benefits of a phased withdrawal retirement income product—as compared to, say, a guaranteed or term life annuity—is the ability to withdraw lump sums at any point within the year, above and beyond the nominated pension payments. We call these drawdowns 'adhoc'.

3.2 Data Preparation

Several super funds provided data at the level of granularity required to support all three components of this methodology. Strategic Insight collected and cleaned the data as part of an ongoing survey initiated by the Institute of Actuaries of Australia. Subsequently, we combined the data to produce the aggregate dataset for analysis. We intend the sample used in this project to be representative of the population of Australian retirees holding phased withdrawal accounts in APRA-regulated superannuation funds.

The dataset analysed was extracted from the available data by taking panel members from two 'entry' cohorts: those observed from the financial year ended 30 June 2004; and those commencing accounts in financial years 2009–11, inclusive. The former represents the earliest available data provided by this fund, while the latter contains data for the newest 'type' of account—those opened on or after 20 September 2007. The first complete financial year observed for these new accounts commenced 1 July 2008 and ended 30 June 2009.

Table 3.1: Minimum Drawdown Rates – Effective Since 1 July 2007

Age	<65	65–74	75–79	80–84	85–89	90–94	95+
Minimum Drawdown Rate	0.04	0.05	0.06	0.07	0.09	0.11	0.14

As the cluster analysis relied heavily on observing as many drawdowns as possible over time for each individual, and since the observation period ended in financial year 2015 for all remaining members, aligning the first years of observation for members within each of the two account types maximised the number of time periods available to compare and contrast individual drawdown behaviours.

As individuals are free to transfer superannuation assets between competing funds, sample exit could occur for at least three reasons:

1. Death
2. Complete withdrawal of account balance
3. Transfer to another fund

We dealt with the first reason for exit by removing all individuals who died while under observation. We assumed that proximity to death has the potential to influence drawdown behaviours, and preferred to focus on the behaviours of surviving retirees in this study. The data for those dying in sample exists in a separate dataset for future analysis.

Secondly, a complete withdrawal of account balance was a behaviour of key interest, and so these individuals remained in-sample. We also retained retirees transferring to another fund in the sample, and study their behaviours while observed.

Overall, this resulted in a sample size of approximately $N = 44,000$ individuals, each observed for $T \in 1, 2, ..., 12$ years.

With account data provided in monthly records, a 'risk appetite' metric could be computed. Comparison of the monthly account investment returns with the S&P/ASX200 index showed extremely high correlation, confirming that investments in Australian equity drove a significant portion of account balance movements—or that where individuals had investments in other markets, these had a high correlation with Australian equity returns. In addition, retirees can customise their investment allocations, varying the proportions they hold in safe and risky assets away from fund defaults. We defined risk appetite as the magnitude of the average ratio between investment returns and corresponding index returns:

$$\text{Risk appetite} = \left| \text{average} \left(\frac{\text{Monthly investment returns in account}}{\text{Corresponding monthly S\&P/ASX200 index return}} \right) \right| \tag{3.2}$$

In this way, high risk appetites correlate with larger equity exposures, while low risk appetites represent smaller returns over time from less variable assets.

For illustrative purposes, Figure 3.1 plots a sample of individual investment return series against the S&P/ASX200 in red.

Other key data manipulations included aggregating the monthly data to 12-month periods corresponding to complete financial years, and transforming variables of interest using the natural logarithm (log) function for modelling purposes.

Figure 3.1: Monthly Investment Returns

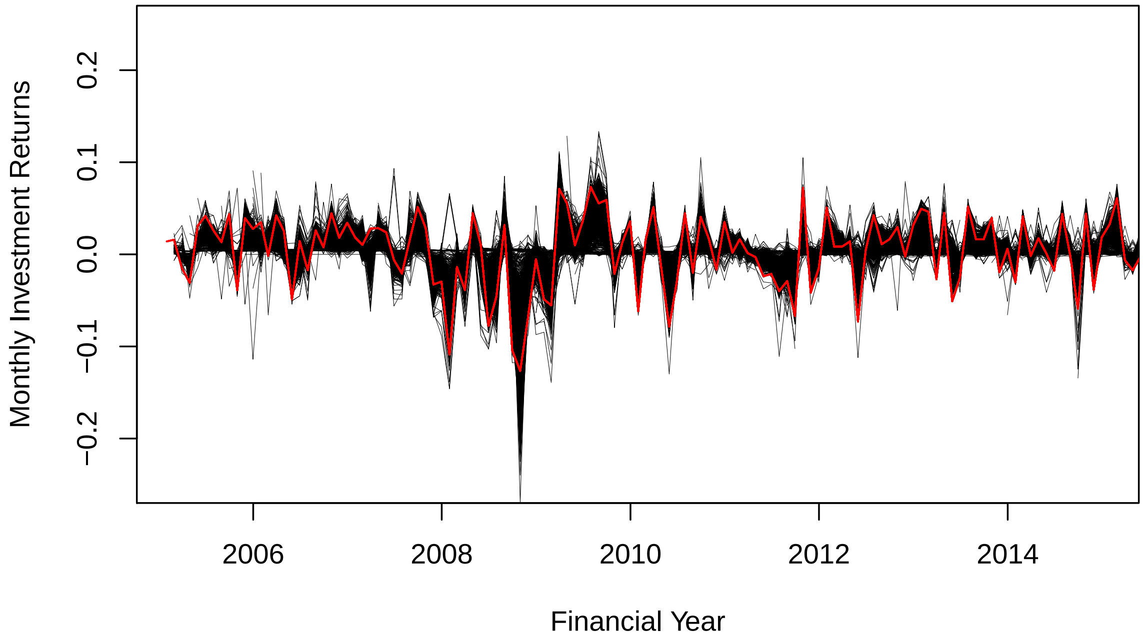

3.3 Component 1: Panel Regression Models

The aim of this component of the methodology is to estimate the effect of the available regressors on an individual retiree's propensity to draw down from their account-based pension. Regression models achieve this by estimating coefficients for each included regressor, and observing their signs, magnitudes and statistical significance. The economic interpretations of interest are, for example, whether drawdown rates or decisions are significantly influenced by the available regressors—such as age, gender and account balance. In particular, regression analysis reports on the effects of regressors after controlling for the values of other included variables. This can disentangle the effects of regressors that are mildly correlated with each other and influence the dependent variable of interest.

Certain regression models can utilise the additional information inherent in data that observes panels—in our case, individual retirees—over time (Wooldridge, 2012, p449). Common panel regression specifications are the Pooled Cross-sectional (PC) regression model, the Fixed Effects (FE) model, and the Random Effects (RE) model.

For illustrative purposes, consider some dependent variable y_{it} which relates to an individual i and is observed through a time index $t = 1, 2, ..., T_i$. This variable could be, for example, the rate at which the individual draws down their account balance, annually. We observe some time-invariant regressors $\mathbf{z}_i$, and other, time-varying, $\mathbf{x}_{it}$. The notation $\mathbf{z}_i$ and $\mathbf{x}_{it}$ represent the column vectors $(z_{i,1}, z_{i,2}, ..., z_{i,K})'$ and $(x_{it,1}, x_{it,2}, ..., x_{it,L})'$, respectively, corresponding to K observed time-invariant characteristics and L time-varying.

3.3.1 Linear Models

In a linear model, we formulate the equation:

$$y_{it} = c + \alpha_i + \mathbf{z}_i'\boldsymbol{\gamma} + \mathbf{x}_{it}'\boldsymbol{\beta} + e_{it} \tag{3.3}$$

Here, α_i represents the unobserved, individual-specific, time-invariant heterogeneity. This expression describes a base level of the dependent variable for individual i, attributed entirely to factors that we cannot observe. Crucially, as indicated by the subscripted i, this effect is, in general, not the same across different individuals. The visible regressors, then, influence how the observed dependent variable fluctuates from this baseline level—with respect to changes in the observed $\mathbf{z}_i$ and $\mathbf{x}_{it}$.

The coefficient vectors $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ combine with the respective class of observed variables to represent linear combinations. That is, $\mathbf{z}_i'\boldsymbol{\gamma} = (\gamma_{i,1}z_{i,1} + \gamma_{i,2}z_{i,2} + ... + \gamma_{i,K}z_{i,K})$ and $\mathbf{x}_{it}'\boldsymbol{\beta} = (\beta_{i,1}x_{it,1} + \beta_{i,2}x_{it,2} + ... + \beta_{i,L}x_{it,L})$. c represents a universal intercept term in the model. Finally, the error term e_{it} absorbs all other unobserved determinants of the dependent variable. Without loss of generality, the error term has mean zero for some value of the constant term c.

Generally, regression models seek to find coefficient estimates which approximate the conditional expectation of the dependent variable given regressor values:

$$E[y_{it}|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i] = c + \alpha_i + \mathbf{z}_i'\boldsymbol{\gamma} + \mathbf{x}_{it}'\boldsymbol{\beta} \tag{3.4}$$

PC models assume that all observations are independent across the time dimension—even successive observations on individual panel members. This assumption is only valid if the regressors capture all individual-specific factors that guide an individual towards some base level of drawdown. That is, the PC model assumes $\alpha_i = 0$. Due to the small set of available regressors, it would be imprudent to assume the data satisfies this condition.

In this respect, FE and RE regressions are more conservative. These models assume that the α_i are nonzero, and therefore induce autocorrelation of the drawdowns made by one individual through time. However, the FE and RE models differ in their treatment of this effect.

FE models remove α_i—and, unfortunately, the $\mathbf{z}_i$—algebraically through the 'within' transformation. First, the mean value of each covariate over the observation period is calculated as $\bar{x}_{i,k} = \frac{1}{T_i}\sum_{t=1}^{T_i} x_{it,k}$. Collectively, in vector notation, $\bar{\mathbf{x}}_i = \frac{1}{T_i}\sum_{t=1}^{T_i} \mathbf{x}_{it}$. A similar calculation is completed for $\bar{y}_i$ and $\bar{e}_i$. Subsequently, the transformed series is given by:

$$\dot{y}_{it} := (y_{it} - \bar{y}_i) = (c - c) + (\alpha_i - \alpha_i) + (\mathbf{z}_i - \mathbf{z}_i)'\boldsymbol{\gamma} + (\mathbf{x}_{it} - \bar{\mathbf{x}}_i)'\boldsymbol{\beta} + (e_{it} - \bar{e}_i) = \dot{\mathbf{x}}_{it}'\boldsymbol{\beta} + \dot{e}_{it} \tag{3.5}$$

More familiar terms for this procedure may be 'de-meaning', or 'centering'. Note that the $\boldsymbol{\beta}$

in equation 3.5 are exactly those from equation 3.4, which we initially intended to estimate. After transformation, the coefficient vector $\boldsymbol{\beta}$ can be obtained by Ordinary Least Squares (OLS) regression. The software program `Stata` can derive the FE estimates of the vector $\boldsymbol{\beta}$ using the `xtreg` command with the `fe` option.

Crucially, the coefficient estimates in a FE model are consistent (asymptotically correct). In our—very—large sample, we will rely on asymptotics in assuming that the FE estimates are the correct values.

In contrast to using FE, RE models can estimate coefficients on the $\mathbf{z}_i$, allowing inference on observed, time-invariant characteristics. However, researchers must take care in assessing the validity of RE models before interpreting the coefficient estimates—particularly due to the strict RE model assumptions, which require zero correlation between α_i and the regressors $\mathbf{z}_i$ and $\mathbf{x}_{it}$.

One can show through matrix algebra (see e.g. Wooldridge, 2010b, ch10) that fitting an RE model can be achieved using OLS estimation of a 'quasi-demeaned' equation:

$$(y_{it} - \theta_i \bar{y}_i) = (c - \theta_i c) + (\alpha_i - \theta_i \alpha_i) + (\mathbf{z}_i - \theta_i \mathbf{z}_i)'\boldsymbol{\gamma} + (\mathbf{x}_{it} - \theta_i \bar{\mathbf{x}}_i)'\boldsymbol{\beta} + (e_{it} - \theta_i \bar{e}_i) \tag{3.6}$$

In this equation, θ_i is given by:

$$\theta_i = 1 - \sqrt{\frac{\sigma_e^2}{T_i \sigma_\alpha^2 + \sigma_e^2}} \tag{3.7}$$

where σ_e^2 and σ_α^2 represent the variance of the random variables e_{it} and α_i in equation 3.4, respectively. `Stata` estimates σ_e^2 and σ_α^2 for unbalanced panels using the methodology of Swamy and Arora (1972).

Since the FE estimates for the coefficients of the $\mathbf{x}_{it}$ are consistent, to trust the RE model results it must prove capable of obtaining the same—or at least, statistically indistinguishable—coefficient estimates on these time-varying regressors. Thus, a crude way to evaluate whether the RE model is valid is to merely inspect how close in value the coefficients are on the $\mathbf{x}_{it}$. The Hausman specification test, however, formalises this comparison.

The Hausman test—implemented in software programs such as `Stata`—aggregates the differences in the coefficients between models, scaled by the relative differences in their precision. Denoting $\boldsymbol{b}$ and $\boldsymbol{B}$ to be the coefficients vectors on the time-varying coefficients derived from the FE and RE models, respectively, we can define:

$$C := (\boldsymbol{b} - \boldsymbol{B})'[(V_{\boldsymbol{b}} - V_{\boldsymbol{B}})^{-1}](\boldsymbol{b} - \boldsymbol{B}) \tag{3.8}$$

where $V_{\boldsymbol{j}}$ represents the variance-covariance matrix for a vector $\boldsymbol{j}$.

This produces a statistic, C, representing the overall dissimilarity between model estimates. Asymptotically, this statistic follows a χ^2 distribution with degrees of freedom equal to the number of compared coefficients, less one. The null hypothesis is that there is no systematic difference in the coefficient estimates between the two models. Rejecting this null indicates statistical evidence that they do differ, and if the FE coefficients are taken to be correct, this implies the RE model is misspecified.

When the data does not support validity of the RE model assumptions, the procedure of Hausman and Taylor (1981) (HT) provides an alternative method for estimating the effect of the time-invariant $\mathbf{z}_i$. By assuming that only some of the $\mathbf{z}_i$ and $\mathbf{x}_{it}$ are uncorrelated with α_i, these 'exogenous' regressors can control for the correlation between the remaining—'endogenous'—regressors and the unobserved effect α_i. Specifically, we may partition the available regressors vectors into exogenous and endogenous components—subscripted 1 and 2, respectively. Thus $\mathbf{x}_{it} = \mathbf{x}_{1it} + \mathbf{x}_{2it}$ and $\mathbf{z}_i = \mathbf{z}_{1i} + \mathbf{z}_{2i}$.

`Stata` implements the HT procedure from Hausman and Taylor (Ibid) as follows. As usual, the aim is to estimate parameters in the model:

$$y_{it} = c + \alpha_i + \mathbf{z}_i'\boldsymbol{\gamma} + \mathbf{x}_{it}'\boldsymbol{\beta} + e_{it} \tag{3.9}$$

Similarly to RE regression, a quasi-demeaning factor θ_i is defined with form:

$$\theta_i = 1 - \sqrt{\frac{\sigma_e^2}{T_i\sigma_\alpha^2 + \sigma_e^2}} \tag{3.10}$$

Performing quasi-demeaning on equation 3.9:

$$(y_{it} - \theta_i\bar{y}_i) = (c - \theta_i c) + (\alpha_i - \theta_i\alpha_i) + (\mathbf{z}_i - \theta_i\mathbf{z}_i)'\boldsymbol{\gamma} + (\mathbf{x}_{it} - \theta_i\bar{\mathbf{x}}_i)'\boldsymbol{\beta} + (e_{it} - \theta_i\bar{e}_i) \tag{3.11}$$

Or more compactly:

$$\tilde{y}_{it} = \tilde{c} + \tilde{\alpha}_i + \tilde{\mathbf{z}}_i'\boldsymbol{\gamma} + \tilde{\mathbf{x}}_{it}'\boldsymbol{\beta} + \tilde{e}_{it} \tag{3.12}$$

To estimate this equation, `Stata` uses instrumental variable (IV) regression of the transformed $\tilde{y}_{it}$ on transformed $\tilde{\mathbf{z}}_i$ and $\tilde{\mathbf{x}}_{it}$. The instruments are exogenous variables $\dot{\mathbf{x}}_{it}$, $\bar{\mathbf{x}}_{1i}$ and $\mathbf{z}_i$—where $\dot{\mathbf{x}}_{it} = (\mathbf{x}_{it} - \bar{\mathbf{x}}_i)$ and $\bar{\mathbf{x}}_{1i} = \frac{1}{T_i}\sum_{t=1}^{T_i}\mathbf{x}_{1it}$. For a detailed review of instrumental variable regression, see for example Chapter 15 of Wooldridge (2012).

Similar to the validation of a RE model, a Hausman specification test can determine the suitability of the HT procedure. The test is conducted as before, but with the HT estimates used in place of the RE estimates.

Since the FE coefficient estimates are asymptotically correct, by comparing the relative signs, magnitudes and statistical significance of the coefficients on the time-varying regressors $\mathbf{x}_{it}$ amongst PC, FE, RE and HT models, we can infer which models fail to satisfy requisite assumptions.

3.3.2 'Nonlinear' Models

While regression equations of the form given in equation 3.4 are often suitable for modelling continuous dependent variables, a class of 'nonlinear' models are more appropriate for modelling non-continuous outcomes—for example, dependent variables taking discrete outcomes. Specifically, we will fit models on binary choice response variables and censored dependent variables.

Binary Choice Models

In binary dependent variable models, the observed response y_{it} is a choice. For example, in time period t an individual i may decide to draw at the minimum drawdown rate (encoded $y_{it} = 1$), or not ($y_{it} = 0$). Here, the conditional expectation of the variable y_{it}, given the values of the observed and unobserved characteristics, is identical to the probability of observing a response ($y_{it} = 1$):

$$E[y_{it}|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i] = 1 \times \Pr(y_{it} = 1|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i) + 0 \times \Pr(y_{it} = 0|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i) = \Pr(y_{it} = 1|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i) \tag{3.13}$$

We refer to the model as nonlinear because we estimate the predicted probability of a response as some general function F applied to a linear combination:

$$\hat{\Pr}(y_{it} = 1|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i) = F(c + \alpha_i + \mathbf{z}_i'\boldsymbol{\gamma} + \mathbf{x}_{it}'\boldsymbol{\beta}) \tag{3.14}$$

One possible choice for the function F is the logistic—inverse logit—function:

$$F(.) = \text{logit}^{-1}(.) = \frac{\exp(.)}{1 + \exp(.)} \tag{3.15}$$

The logistic function transforms a variable on $(-\infty, \infty)$ to $(0, 1)$—making it suitable for translating an unrestricted linear combination into a meaningful probability value.

Coefficients in a logistic regression model are interpreted as changes in the log odds ratio—relative changes in the odds ratio—due to unit changes in the regressors.

Unfortunately, in nonlinear models, the within transformation used in linear FE models can no longer remove the unobserved α_i algebraically. RE methods can be extended to nonlinear models, however these inherit the main constraint of linear RE models: the strict assumptions require that the unobserved α_i is uncorrelated with the observed regressors $\mathbf{z}_i$ and $\mathbf{x}_{it}$. Furthermore, although it would be straightforward to estimate nonlinear models using PC, these models are misspecified whenever successive observations for a panel member are not independent over time.

To avoid both the PC model and the strong RE assumptions, we will use Correlated Random Effects (CRE) models—which use techniques to control for the potential correlation between the available regressors and the α_i. Wooldridge (2010a) attributes the CRE model in balanced panels to Chamberlain (1982) as a revision to the work of Mundlak (1978). Wooldridge also extends nonlinear CRE models to unbalanced panels.

CRE can be related to RE as follows. In the general nonlinear case, we have the regression

model:
$$E[y_{it}|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i] = F(c + \alpha_i + \mathbf{z}_i'\boldsymbol{\gamma} + \mathbf{x}_{it}'\boldsymbol{\beta}) \tag{3.16}$$

While RE would assume that the unobserved α_i follows some distribution—for example, Gaussian—with mean 0 and variance σ_α^2, CRE uses time-invariant information to control for any correlation between the α_i and the regressors $\mathbf{z}_i$ and $\mathbf{x}_{it}$. For the $\mathbf{x}_{it}$, which are time-varying, CRE uses the time-averaged level $\bar{\mathbf{x}}_i = \frac{1}{T_i}\sum_{t=1}^{T_i}\mathbf{x}_{it}$ to control for the correlation with α_i. Therefore, in a CRE model, the original α_i has been replaced with a new α_i with mean $\mathbf{z}_i'\boldsymbol{\gamma} + \bar{\mathbf{x}}_i'\boldsymbol{\delta}$ and variance σ_α^2. As a result, the $\mathbf{z}_i'\gamma$ terms in equation 3.16 move from being explanatory variables for the dependent variable y_{it} to being controls for the unobserved heterogeneity.

The new conditional expectation of the response variable becomes:

$$E[y_{it}|\mathbf{x}_{it}, \mathbf{z}_i, \alpha_i] = F(c + (\mathbf{z}_i'\boldsymbol{\gamma} + \bar{\mathbf{x}}_i'\boldsymbol{\delta}) + \mathbf{x}_{it}'\boldsymbol{\beta}) \tag{3.17}$$

with both $\mathbf{z}_i$ and $\bar{\mathbf{x}}_i$ being used to control for the level of α_i. For implementation purposes, this is identical to running a RE model—with the inclusion of the new time-averaged $\bar{\mathbf{x}}_i$ as regressors. `Stata` implements estimation of the RE Logit model via the command `xtlogit`.

Crucially, the $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ coefficients estimated by CRE models do not have the desirable interpretation as partial effects on the response variable (Wooldridge, 2010a). Instead, only the $\boldsymbol{\beta}$ coefficients have the usual interpretation. This is one reason why we do not implement the CRE method for linear models: the $\boldsymbol{\beta}$ coefficients are readily available using the simpler FE estimation method.

Censored Regression Models

Dependent variables may also be censored, due either to limitations in data collection, or natural constraints on the range of a dependent variable. In either case, the presence of probability masses at certain values of the dependent variable distribution causes regular OLS techniques—which assume a continuous dependent variable with unrestricted support—to produce biased coefficients, due to the limited range of the dependent variable.

For illustrative purposes, assume an observed response variable y_{it} is censored from above and below by the values b and a, respectively. That is, the observed variable appears to be continuous on the interval (a, b), but contains significant probability masses at both a and b. Of economic interest is how changes in the values of the regressors $\mathbf{z}_i$ and $\mathbf{x}_{it}$ influence changes in the response y_{it}—which only has a meaningful interpretation for the continuous portion of the distribution.

To avoid biased OLS estimates in this scenario, we can specify a latent (underlying) variable y_{it}^* which is not censored:
$$y_{it}^* = c + \alpha_i + \mathbf{z}_i'\boldsymbol{\gamma} + \mathbf{x}_{it}'\boldsymbol{\beta} + e_{it} \tag{3.18}$$

What we observe instead is the censored version of this true, underlying behaviour:

$$y_{it} = \begin{cases} a, & y^*_{it} \leq a \\ y^*_{it} = c + \alpha_i + \mathbf{z}_i'\boldsymbol{\gamma} + \mathbf{x}_{it}'\boldsymbol{\beta} + e_{it}, & a < y^*_{it} \leq b \\ b, & y^*_{it} > b \end{cases} \tag{3.19}$$

By estimating the coefficients of the latent variable model in equation 3.19, we obtain the desired partial effects.

With cross-sectional data, tobit models can estimate coefficients in situations where an otherwise continuous variable has significant probability masses at one or both edges of its support. For panel data, `Stata` implements tobit models through the command `xttobit`. In this case, censoring also prevents a within transformation from removing the unobserved α_i. Instead, we must rely on RE model estimation methods. However, as in the binary choice model, CRE models can correct for the correlation between the unobserved α_i and the regressors $\mathbf{z}_i$ and $\mathbf{x}_{it}$ through inclusion of $\bar{\mathbf{x}}_i$ as an additional regressor. Similar limitations on the interpretability of coefficients apply, as described for the binary choice models.

3.3.3 Model Validation

In addition to drawing statistical inference from regression output tables, we are interested in how much of the overall variability of the observed responses can be explained by the available regressors. For the linear and censored regression models, we will inspect residual diagnostics. In general, a model that fits the data well will have no discernible pattern in the residuals—both on aggregate and when plotted against the fitted values and individual regressors.

For the binary choice models, however, where the observed values are binary but the predicted values take a range of probabilities, the residuals are less meaningful. Instead, we will inspect the ability of the model to classify individuals—broadly, how often the model is correct when predicting a response or no response.

3.3.4 Regressors and Regressands

The panel models study five dependent variables of interest:

1. Decision to draw at the minimum rate in a given financial year, or not
2. Decision to make an adhoc drawdown in a given financial year, or not
3. The excess regular drawdown rate over a financial year, conditional on having drawn above the minimum
4. The unconditional regular drawdown rate over a financial year
5. The adhoc drawdown rate over a financial year, conditional on having made an adhoc drawdown

For each of these models, we consider the following list of available—or constructed—variables as candidate regressors, categorised as either time-varying (TV) or time-invariant (TI).

- Age at financial year start (TV)
- Account balance at financial year start (TV)
- The minimum drawdown rate—for that member in that financial year (TV)
- Financial-year dummy variables (TV)
- Gender (TI)
- Age at account open—a proxy for retirement age (TI)
- Risk appetite (TI)
- Age at 31 December 2015—the cohort effect (TI)

When including a set of dummy variables in standard, cross-sectional regression models, multicollinearity is avoided by dropping one variable in the set. In panel models, however, some situations require dropping more than one time dummy. This is because any variable that increases by one in each successive observation of an individual is indistinguishable from the passage of time (measured in years). If only one time dummy was dropped and we had one or more of these unit-incrementing variables, the multicollinearity issue would resurface.

In general, for each variable we include that increases by one between subsequent time indices—in our case, the 'age at financial year start' variable—we must drop one additional time dummy variable. As a result, although our complete set of financial year dummy variables covers 2004 to 2015 inclusive, we must drop two in our regression modelling, and this pair of years becomes the 'base case' against which we can compare the effect of the remaining years. As the earlier years in our sample exhibit more interesting effects than later years, we select 2014 and 2015 to be the base case, and include time dummy variables for the 2004 to 2013 financial years, inclusive.

The age definition is the age at the start of the relevant financial year, to reflect the rules in the legislation for determining which minimum drawdown rate applies to the individual during a particular financial year.

When modelling dependent variables using linear models, we will first transform nonzero drawdown rates—naturally constrained on (0, 1]—using the natural logarithm (log) function. This spreads out the support of the distribution, reduces skewness and increases symmetry—three changes which make the dependent variable more suitable to modelling by the techniques described in this section.

3.4 Component 2: Cluster Analysis and Identification of Behavioural Groups

The aim of this component is to allocate individuals in the sample into behavioural groups based on their observed drawdowns.

3.4.1 Manual Grouping

The presence of minimum drawdown rates as a default option immediately suggests a potential behavioural group. So too does the optimality literature, which suggests that drawing at level rates or level amounts might be heuristics retirees employ in decision-making. Thus, manual identification of some behavioural groups may be possible using filters on the observed data. For example, a rule which finds individuals drawing at or near the minimum drawdown rates for most observed periods would identify the members of one of these behavioural groups.

Specifically, we search for the following five groups in the data:

1. Draw at—or very near—the minimum drawdown rates in all—or most—observed periods
2. Similar to group 1, although do not adjust to the concessional minimum rates applying for financial years 2009–13 inclusive
3. Draw at—or very near—10% of their account balance annually. In Transition to Retirement Income Products (TRIPs), this is the maximum allowable rate of account drawdown
4. Draw the same dollar amount from their accounts in all—or most—observed periods
5. Draw at the same rate from their accounts in all—or most—observed periods, exclusive of individuals in groups 1, 2 or 3

3.4.2 Machine-Assisted Grouping

Where imagination and energy limit the extent of classification by a manual grouping method, a machine-assisted extension can add further value. Cluster analysis can create groups of individuals using observed characteristics (James et al., 2013), and of interest to the research is grouping individuals by their drawdown behaviours across the panel. This can be achieved by treating each annual drawdown rate as a separate variable ('characteristic') for the individual, and grouping based on the observed drawdown rates over time.

The problem can be visually expressed using the toy example in Figure 3.2. Drawdown rate, as a proportion of account balance, is given on the vertical axis, while the horizontal axis tracks each individual over time.

A successful cluster analysis would find four 'clusters' in the toy dataset. Individuals 1, 2, and 5 inhabit their own cluster, while individuals 3 and 4 comprise the fourth.

As well as clustering individuals based on patterns in the level of their drawdowns, it may prove instructive to remove the impact of the starting level. If the first differences are taken in the series, the 'shape' of the drawdown pattern forms the basis for clustering, rather than the actual dollar amount or proportion of account balance drawn. For illustrative purposes, consider Figure 3.3, where the underlying toy dataset is the same, but the series of interest is the first difference in the drawdown rates.

Figure 3.2: Toy Example for Cluster Analysis

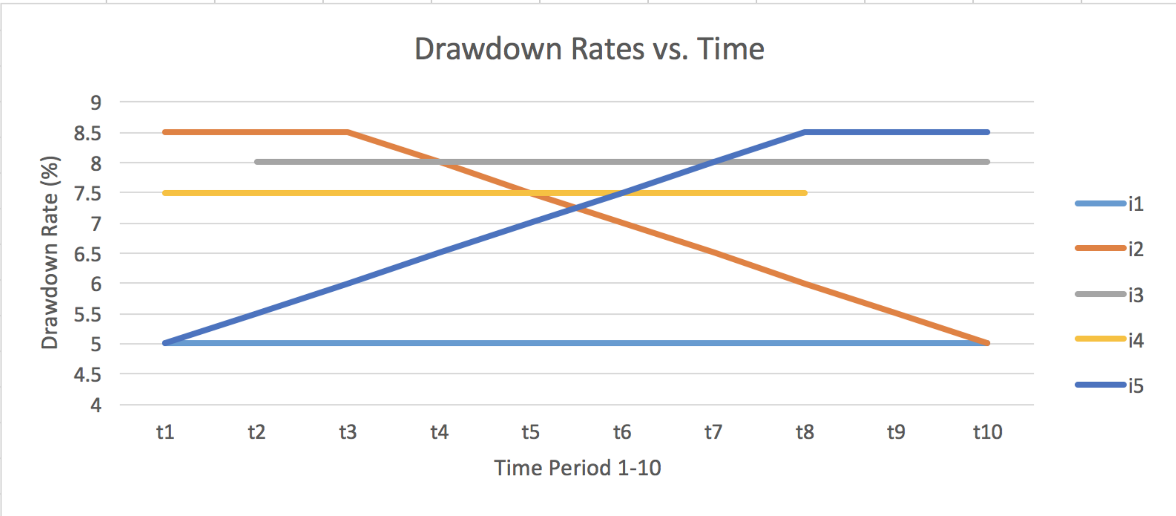

individual	t1	t2	t3	t4	t5	t6	t7	t8	t9	t10
i1	5	5	5	5	5	5	5	5	5	5
i2	8.5	8.5	8.5	8	7.5	7	6.5	6	5.5	5
i3		8	8	8	8	8	8	8	8	8
i4	7.5	7.5	7.5	7.5	7.5	7.5	7.5	7.5		
i5	5	5.5	6	6.5	7	7.5	8	8.5	8.5	8.5

Figure 3.3: Toy Example for Cluster Analysis – First Differenced

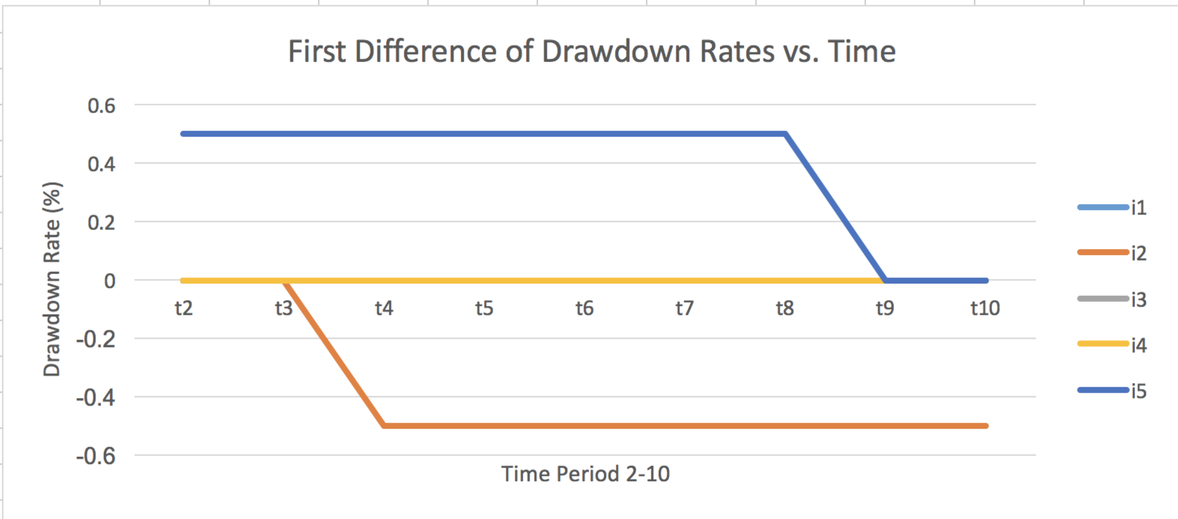

individual	t1	t2	t3	t4	t5	t6	t7	t8	t9	t10
i1	NA	0	0	0	0	0	0	0	0	0
i2	NA	0	0	-0.5	-0.5	-0.5	-0.5	-0.5	-0.5	-0.5
i3	NA	NA	0	0	0	0	0	0	0	0
i4	NA	0	0	0	0	0	0	0	NA	NA
i5	NA	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0	0

In this example, the three individuals drawing constant rates in Figure 3.2 would be grouped into one cluster, while the increasing and decreasing drawdowns would become two additional clusters.

In solving problems of this nature, two cluster analysis algorithms are common: k-means and hierarchical clustering (James et al., 2013). The k-means method requires a distance metric to be calculated across all applicable rows and columns. However, in the available dataset, not all accounts are observed for the entire sample duration, creating missing data. While the issue of missing data in k-means cluster analysis has some—arguably suboptimal—solutions, involving deletion or imputation of data, new algorithms like the `k-POD` `R` package present other compelling options for k-means analysis in the presence of missing data (Chi et al., 2016).

Alternatively, hierarchical clustering proceeds unhindered in the presence of missing data. In this procedure, a dissimilarity measure compares different individuals, while a linkage criterion allows computation of dissimilarities between clusters—which may be comprised of more than one individual. Crucially, the choice of linkage method can significantly impact the computational time of hierarchical cluster analysis as data sets grow in size (Murtagh, 1983). In the current methodology, we experiment with a variety of distance metrics and linkage methods. By visually inspecting the clustering results under different combinations, we find the combination which maximises within-cluster similarity and between-cluster dissimilarity. As the behaviours we seek require an economic interpretation, through inspection we naturally find the clustering parameters which provide the most meaningful behavioural results.

Consequently, we perform hierarchical clustering using the statistical software program `R` and the `cluster` package. Panel visualisations of individual drawdown trajectories through time will both motivate the exploration of particular clusters, and confirm sensible clustering results. Furthermore, as outlined in Chapter 2, we intend to compare the obtained clusters with the drawdown strategies suggested by Horneff et al. (2008) and Bateman and Thorp (2008).

3.5 Component 3: Categorical Regression for Behavioural Group Allocation

After allocating individuals into groups corresponding to their drawdown behaviours, the role of a categorical regression model is to comment on the statistically significant differences in the characteristics of retirees displaying disparate drawdown behaviours. Moreover, we are interested in how the available characteristics—including age, gender and account balance—determine the relative and absolute probabilities of adhering to one of the identified behaviours.

Specifically, individual i belonging to cluster j is denoted $C_i = j$. The probabilities estimated by the model will have the form:

$$\Pr(C_i = j|\mathbf{z}_i) = \theta_{ij}(\mathbf{z}_i), \quad j = 1, 2, ..., J \tag{3.20}$$

where $\mathbf{z}_i$ is a vector of length $k+1$ consisting of the constant 1 followed by k relevant explanatory variables available for individual i, i.e. $\mathbf{z}_i = (1, x_{i,1}, x_{i,2}, ..., x_{i,k})'$.

Figure 3.4 illustrates the intended outputs from the categorical regression model.

One method to construct the function θ_{ij} is by extending the functional form of a logistic binary choice regression to the multinomial case:

$$\theta_{ij}(\mathbf{z}_i) = \frac{\exp(\mathbf{z}_i'\boldsymbol{\beta}_j)}{\sum_{m=1}^{J}\exp(\mathbf{z}_i'\boldsymbol{\beta}_m)} \tag{3.21}$$

Here $\boldsymbol{\beta}_j$ is a vector of $k+1$ coefficients, and $\mathbf{z}_i'\boldsymbol{\beta}_j$ represents the linear combination $\beta_{j,0} + \beta_{j,1}z_{i,1} + \beta_{j,2}z_{i,2} + ... + \beta_{j,k}z_{i,k}$.

This model then uses the available regressors to predict the probability θ_{ij} of an individual i belonging to any one of the J clusters, conditional on the observed values of the regressors $\mathbf{z}_i$. Parameter estimation is achieved by Maximum Likelihood Estimation (MLE), where the joint likelihood L of the observed data is given by:

$$L = \prod_{i=1}^{N}\prod_{j=1}^{J}\{\theta_{ij}(\mathbf{z}_i)\}^{1_{\{C_i=j\}}} = \prod_{i=1}^{N}\prod_{j=1}^{J}\left\{\frac{\exp(\mathbf{z}_i'\boldsymbol{\beta}_j)}{\sum_{m=1}^{J}\exp(\mathbf{z}_i'\boldsymbol{\beta}_m)}\right\}^{1_{\{C_i=j\}}} \tag{3.22}$$

One crucial property of the θ_{ij} function as specified in equation 3.21 is the adherence to the law of total probability, $\sum_{j=1}^{J}\theta_{ij} = 1$. For this reason, one set of coefficients $\boldsymbol{\beta}_j$ are redundant—one group can be selected as the base case, and the corresponding predicted probability is known as $\theta_{i,\text{base}} = 1 - \sum_{j\neq\text{base}}\theta_{ij}$. Coefficients for other groups are then interpreted as relative changes in the log odds ratio of belonging to a particular cluster—relative to the base group.

We use `Stata` to fit the model, providing coefficient estimates and standard errors.

This specification of θ_{ij}, which is based on the logistic function in the binary variable case, is not guaranteed to provide the best—or even a good—fit to the data. Consequently, after fitting the model, validation will uncover how often behaviours can be successfully explained or predicted using the available regressors.

One practical device to assess this fit is the 'confusion' matrix—an extension of the binary choice classification table. The number of individuals correctly predicted by the model to belong to a particular cluster is compared to two key quantities: how many predictions were made for that cluster, both correctly and erroneously; and how many individuals originated from that cluster in total.

A relevant conceptual point regarding the data used in components 1 and 2 deserves further elaboration. The fitting process requires observation of C_i, the group which the individual i was allocated to as a result of the cluster analysis. The cluster analysis therefore reduces the dimensionality of the observed 'behaviours' to 1, allowing this behaviour to become a time-invariant dependent variable for the individual, constructed using the drawdown experience

Figure 3.4: Visual Representation of Categorical Regression Model Prediction

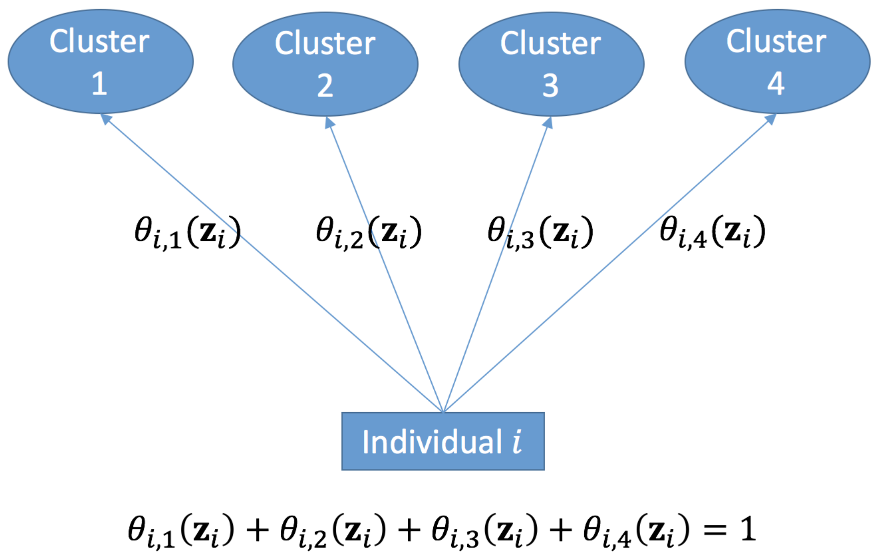

over time. Consequently, in predicting which cluster an individual may belong to over time, only time-invariant variables, known at the beginning of the sample period, will be used as the independent variables for the multinomial regression. Hence the dependent and independent variables available for the multinomial regression are a subset of those listed earlier in Section 3.2 on Data Preparation—plus a time-invariant version of the account balance variable, captured in the first year of observation for each panel member.

- Observed response variable:
  - $C_i = j$, i.e. i belongs to behavioural group j, for $j = 1, 2, ..., J$
- Explanatory variables—time-invariant (RHS):
  - Account balance at first observation
  - Gender
  - Age at account open
  - Risk appetite
  - Age at 31 December 2015
- Estimated/Predicted response variable (LHS):
  - J probabilities for each individual, each corresponding to the probability of belonging to one of the J behavioural groups

3.6 Summary of Methodology

Our approach to analysing the drawdown data has three major components. First, panel regression modelling techniques relate observed drawdown rates to retiree characteristics. These models indicate the direction, magnitude and statistical significance of the effect of the regressors on the dependent variables. Second, a combination of manual grouping and machine-assisted cluster analysis allocates retirees into distinct behavioural groups—characterised by their observed drawdowns over time. Finally, a categorical regression model finds the statistical relationships between available characteristics and the likelihood of observing a specific behaviour within phased withdrawal arrangements.

CHAPTER 4

RESULTS

This chapter presents the results from all three methodology components, and further explores interesting features of the data.

4.1 Component 1: Panel Regression Models

4.1.1 Preliminary Analysis

First, we numerically and visually summarise the dependent variables of interest and the candidate regressors.

Drawdown at the Minimum Rates

Table 4.1 shows the aggregate proportion of drawdowns (44%) in sample which occur at the minimum drawdown rates effective since 1 July 2007. For the subsequent modelling, in years where concessional rates applied below the regular minima, we continue to use the unmodified minimum drawdown rates. In this way, we capture all individuals deciding to draw at the legislated minimum rates, regardless of whether they were aware of the temporary introduction of concessional rates. Additionally, ignoring the minimum rates allows easier comparison of our overall results with previous analytical work done by Plan For Life (2016). For comparison, Table 4.1 also displays the proportion of drawdowns (32%) made at the true effective minima, including the concessional rates applying to financial years ended 30 June 2009–13, inclusive.

Making Adhoc Drawdowns

Table 4.1 shows that 12% of the observed annual drawdowns in our sample contained an adhoc withdrawal.

Regular, Excess and Adhoc Drawdown Rates

Figures 4.1, 4.2 and 4.3 show distributions of the regular drawdown rates, excess regular drawdown rates, and adhoc drawdown rates, respectively. Note that the excess regular drawdown rate is conditional on observing drawdown above the minimum rates. Similarly, the adhoc drawdown rate is conditional on observing a nonzero adhoc drawdown. Note also that we right-censor the adhoc drawdown rate data at 90% of account balance. This is to reduce the noise created by individuals who make regular drawdowns, at or near the minimum rates, and then subsequently draw the remainder of their account balance ad hoc, during the same financial year.

Regressor Properties

In determining which of the candidate regressors may be suitable for inclusion in the modelling procedure, we are cautious of the pairwise correlations between several of our age-related variables. In regression modelling, high collinearity between several regressors reduces the precision with which we can estimate the effects of any one of the correlated set. The combined histogram, pairwise scatterplot and pairwise correlation matrix of Figure 4.4 assists with preventing the high collinearity issue.

We observe that the cohort effect 'Age at 31 December 2015' shows high pairwise correlation with the other age variables, as well as the minimum drawdown rate, which is a function of age. To avoid introducing high collinearity into our regression models, we omit the cohort effect.

The other large correlation statistic is between the time-varying age variable and the time-invariant age at which a member opens their account—a proxy for the retirement age. Modelling results did not suggest that including this variable impacted on the standard errors of the other regressors, and so we retained this variable in our modelling procedure. We also present numerical summary statistics on the candidate regressors in Table 4.2.

Table 4.1: Summary of Observed Binary Choice Outcomes

Behaviour	Observed Frequency
Draw at Minimum Rate (Concessional or Non-Concessional)	0.440
Draw at Minimum Rate (Concessional Only)	0.321
Make Adhoc Drawdown	0.124

Figure 4.1: Histogram of Regular Drawdown Rate

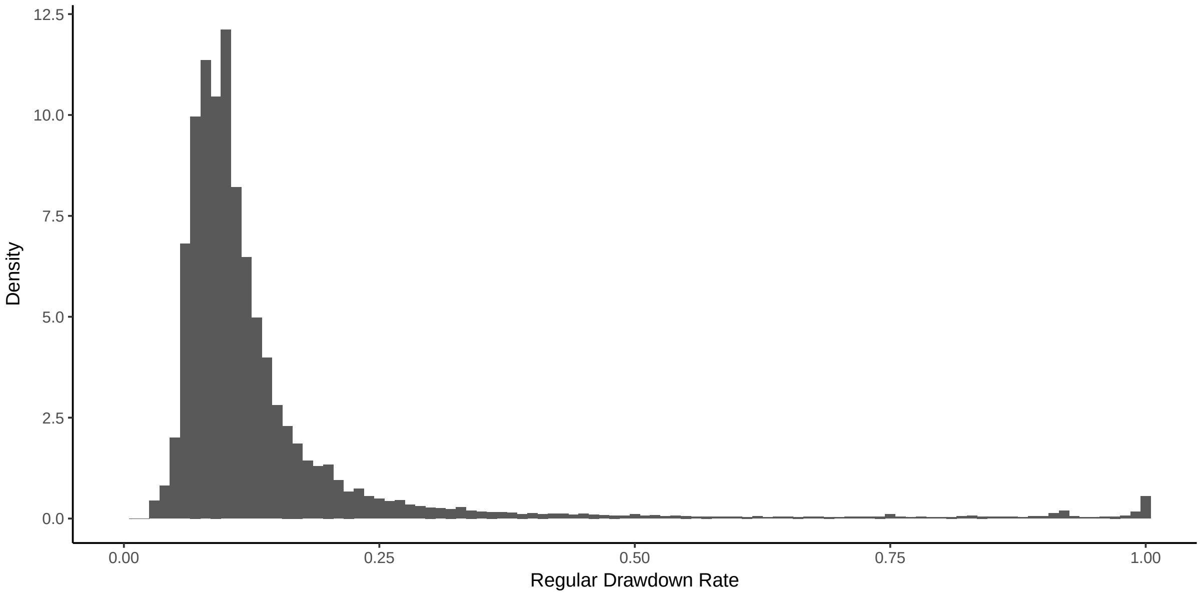

Figure 4.2: Histogram of Excess Regular Drawdown Rate

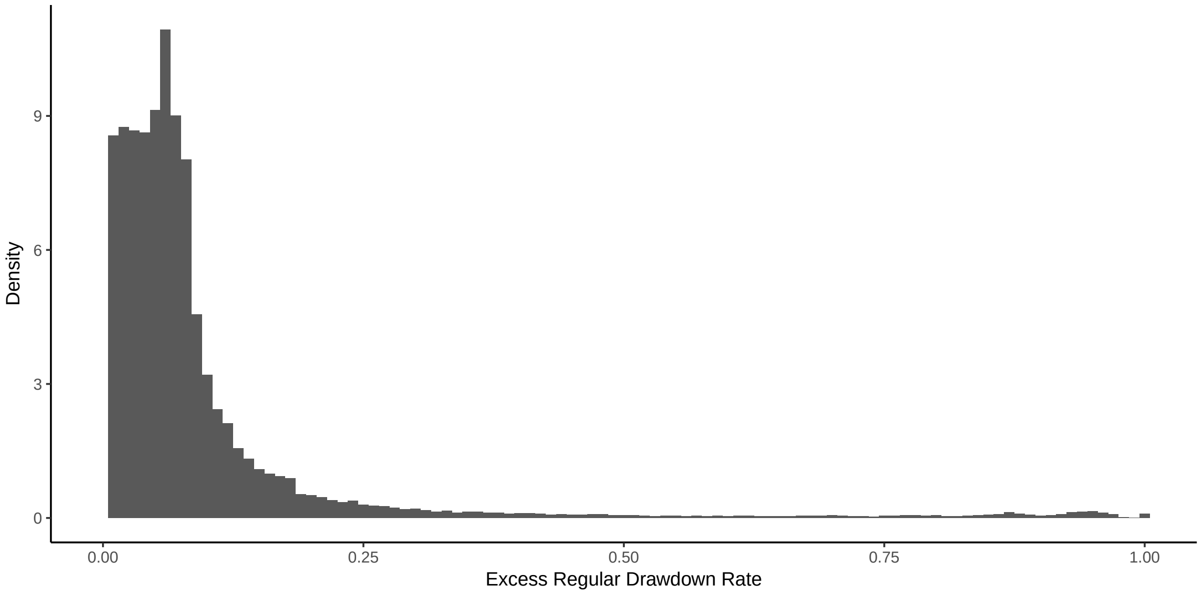

Table 4.2: Summary Statistics for Candidate Regressors – Panel Modelling

Variable	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Age	55	65	69	69.6	74	99
Account Balance	$0	$39,085	$73,086	$116,312	$136,037	$4,952,911
Risk Appetite	0.00	0.30	0.47	0.46	0.61	2.00
Age at Account Open	40.2	60.6	64.1	63.7	66.0	89.4
Age at 31 December 2015	57.7	72.4	77.7	76.94	81.7	103.8
Gender = Male	0	0	1	0.571	1	1
Legacy Account	0	1	1	0.650	1	1

Figure 4.3: Histogram of Adhoc Drawdown Rate

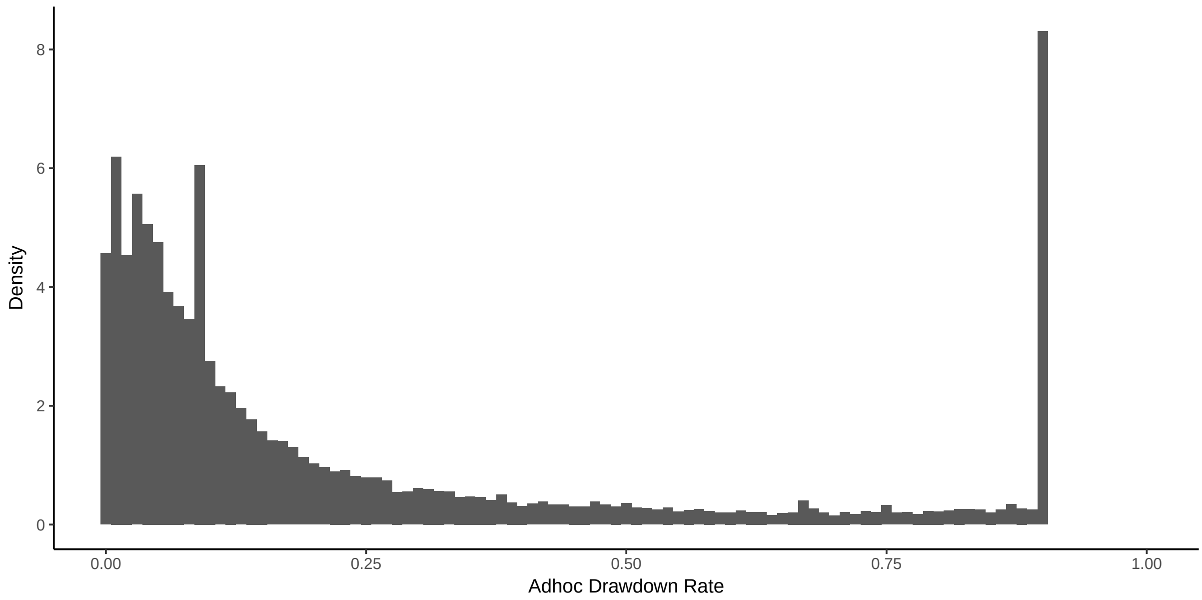

Figure 4.4: Pairwise Correlation Matrix of Candidate Regressors

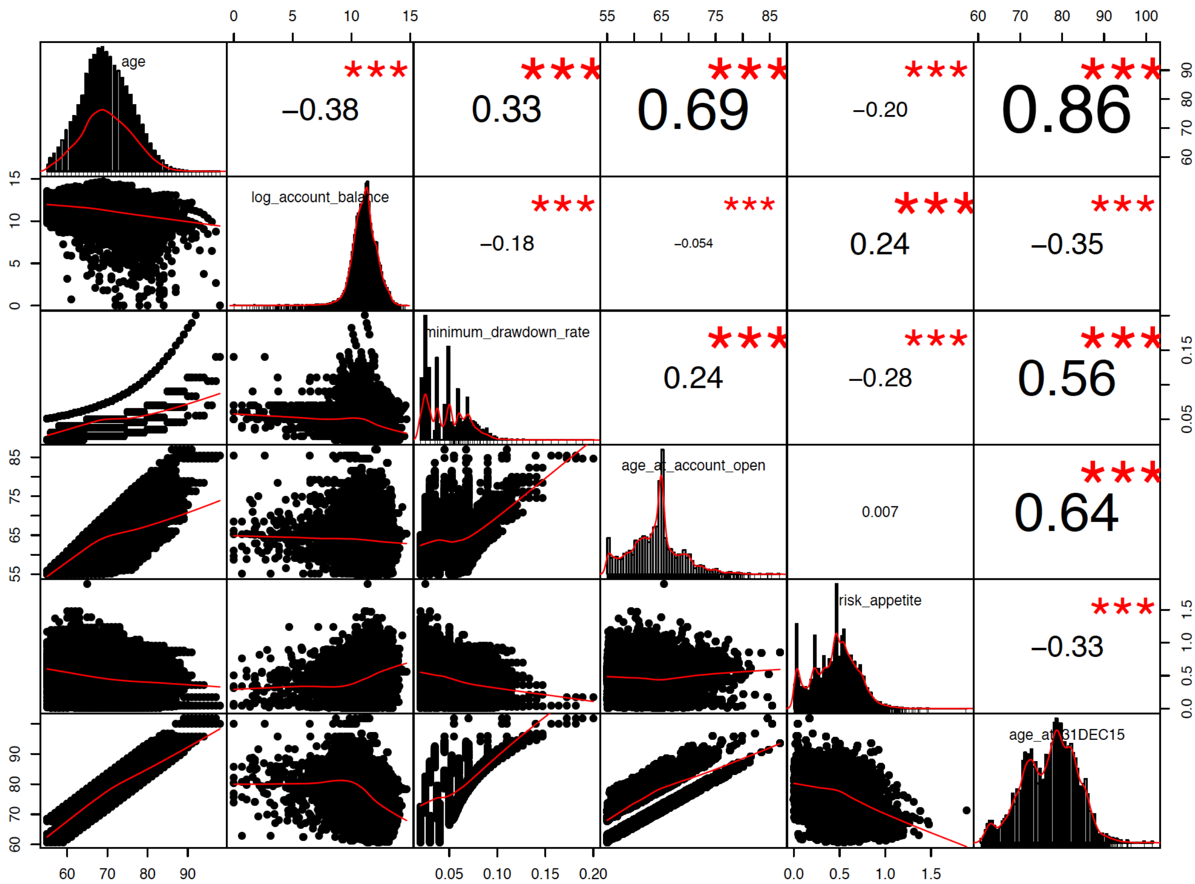

4.1.2 Binary Choice Models

We model two pairs of binary choices, which members of the panel made during each year of observation:

1. Drawing at the minimum rate in a given financial year, or not
2. Making an adhoc drawdown in a given financial year, or not

For both of these pairs of choices, we aim to fit and evaluate a CRE Logit model. For comparison purposes, we also fit models of Pooled Cross-sections and (uncorrected) Random Effects. Although validity of these latter models requires stricter assumptions than our model is likely to meet, for completeness we provide them to show what the impact on our estimated coefficients would have been if we had not applied the more flexible CRE model.

Decision to Draw at Minimum Rate

Table 4.3 contains the Pooled Cross-sectional (PC), Random Effects (RE) and Correlated Random Effects (CRE) Logit model estimates. Using this, we can judge the statistical significance of the available regressors in estimating the probability that an individual draws at the minimum.

In Logistic regression models, regressor coefficients are interpreted as relative changes to the log odds ratio for a corresponding unit change in the regressor value. As these are not directly the changes to the response probability, in Table 4.4 we provide marginal effects that are directly interpretable as the change in probability of observing a response, given a unit change in the corresponding regressor. As the marginal effects in a Logit model depend not only on the value of the regressor of interest, but also on the level of all other variables in the model, we report the average marginal effects (AMEs). We obtain the AMEs using `Stata`'s `margins` command, with the `dydx` option. This approximates the marginal effects of a regressor of interest for each observation in the sample, given the level of the other regressors for that observation, and averages the resulting individual marginal effects across all observations. We interpret the reported marginal effects for several variables of interest.

For the age variable, only the linear effect is statistically significant at least at the 1% level. An incremental year of age implies an average increase of 2.1% in the probability of drawing at the minimum rates.

Although the coefficient on the log account balance is negative, the positive coefficient on its square term begins to dominate very early. Above approximately $9500, increasing the account balance increases the probability of drawing at the minimum. Moving from an account balance of $100,000 to $110,000 increases the probability of drawing at the minimum by approximately 0.6%, while moving from an account balance of $1,000,000 to $1,100,000 increases the same probability by about 1.1%. Thus while statistically significant, the effect on rising account balances is relatively insignificant economically.

Table 4.3: Draw at Minimum Rate – Binary Choice Regression Model Output

	PC Logit Model	RE Logit Model	CRE Logit Model
Age	0.0987***	0.0410	0.278**
	(0.0150)	(0.0813)	(0.105)
Age^2	−0.00103***	−0.00137*	−0.000997
	(0.000110)	(0.000616)	(0.000718)
Log Account Balance	−0.575***	−2.689***	−3.007***
	(0.0232)	(0.0900)	(0.136)
(Log Account Balance)^2	0.0390***	0.151***	0.164***
	(0.00116)	(0.00490)	(0.00791)
Minimum Drawdown Rate	16.37***	42.16***	45.24***
	(1.244)	(5.006)	(5.329)
Financial Year = 2004	0.582***	0.938***	3.984***
	(0.0479)	(0.204)	(0.302)
Financial Year = 2005	0.789***	1.517***	4.272***
	(0.0472)	(0.201)	(0.283)
Financial Year = 2006	0.902***	1.874***	4.355***
	(0.0471)	(0.198)	(0.265)
Financial Year = 2007	0.999***	2.151***	4.341***
	(0.0471)	(0.197)	(0.249)
Financial Year = 2008	−2.252***	−5.138***	−3.168***
	(0.0461)	(0.125)	(0.186)
Financial Year = 2009	−0.842***	−2.038***	−0.273
	(0.0394)	(0.137)	(0.205)
Financial Year = 2010	−0.265***	−0.607***	0.842***
	(0.0370)	(0.131)	(0.184)
Financial Year = 2011	−0.0259	−0.0709	1.054***
	(0.0362)	(0.129)	(0.169)
Financial Year = 2012	−0.0704**	−0.165*	0.624***
	(0.0254)	(0.0721)	(0.0987)
Financial Year = 2013	0.00264	0.000984	0.481***
	(0.0261)	(0.0700)	(0.0840)
Risk Appetite	−0.342***	−1.330***	−1.422***
	(0.0239)	(0.125)	(0.128)
Gender = Male	−0.241***	−0.727***	−0.673***
	(0.0103)	(0.0550)	(0.0540)
Age at Account Open	0.0643***	0.208***	0.189***
	(0.00274)	(0.0135)	(0.0145)
Legacy Account	−0.624***	−1.416***	−1.168***
	(0.0275)	(0.137)	(0.188)
($\bar{\mathbf{x}}_i$ omitted)	.	.	.
Constant	−4.848***	0.860	−47.15***
	(0.519)	(2.754)	(3.791)
Observations	199334	199334	198696

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4.4: Draw at Minimum Rate – Binary Choice Model Average Marginal Effects

	CRE Average Marginal Effects	
Age	0.0210**	(0.00793)
Age^2	−0.0000753	(0.0000542)
Log Account Balance	−0.227***	(0.00990)
(Log Account Balance)^2	0.0124***	(0.000580)
Minimum Drawdown Rate	3.417***	(0.400)
Financial Year = 2004	0.301***	(0.0225)
Financial Year = 2005	0.323***	(0.0210)
Financial Year = 2006	0.329***	(0.0196)
Financial Year = 2007	0.328***	(0.0184)
Financial Year = 2008	−0.239***	(0.0141)
Financial Year = 2009	−0.0206	(0.0155)
Financial Year = 2010	0.0636***	(0.0139)
Financial Year = 2011	0.0796***	(0.0127)
Financial Year = 2012	0.0471***	(0.00742)
Financial Year = 2013	0.0363***	(0.00631)
Risk Appetite	−0.107***	(0.00958)
Gender = Male	−0.0509***	(0.00403)
Age at Account Open	0.0143***	(0.00107)
Legacy Account	−0.0882***	(0.0142)
Observations	198696	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Increasing the minimum drawdown rate by 0.01—for example moving from a minimum drawdown rate of 5% to 6% of account balance—increases the probability of drawing at the minimum by 3.4% on average.

Compared to the base case financial years 2014 and 2015, drawdowns were roughly 30–33% more likely to occur at the minimum in financial years 2004 through 2007. In financial year 2008, drawdowns were about 24% less likely to be at the minimum. In other financial years, the relative probabilities were more modest in magnitude.

In a CRE model, the coefficients of the time-invariant regressors—such as Gender and Age at Account Open—cannot be directly interpreted as impacting the dependent variable of interest. Instead, these regressors act—alongside the average of the time-varying regressors—as controls for the unobserved heterogeneity α_i.

To evaluate how well the CRE Logit model classifies the in-sample responses, we inspect the classification results in Table 4.5. The columns represent the number of data points at which people did draw above the minimum and at the minimum, respectively. The rows indicate the number of observations that the model predicted in-sample to draw above the minimum and at the minimum, respectively. Collectively, this matrix can be used to derive the Sensitivity, Specificity, Positive Predictive Value (PPV), Negative Predictive Value (NPV) and Overall Accuracy provided in Table 4.6. These four quantities are defined as follows:

- Sensitivity = Pr(predict response | observe response)
- Specificity = Pr(predict no response | observe no response)
- Positive Predictive Power = Pr(observe response | predict response)
- Negative Predictive Power = Pr(observe no response | predict no response)

To translate the predicted probabilities that emerge from the model into the binary choice outcome, a cutoff value of 0.5 was used. Predicted probabilities at least as large as 0.5 were classified as a response (drawing at the minimum), and vice-versa for no response (drawing above the minimum). Varying this cutoff value from 0.5, we were unable to find a cutoff point that materially raised the Overall Accuracy of the model.

These in-sample classification results show that our selection of mainly administrative variables could not capture most of the variation in the decisions made, even after using the Correlated Random Effects model to approximate the contribution of each individual's unobserved heterogeneity.

Table 4.5: Draw at Minimum Rate – Binary Choice Classification Table

	Observed		
Predicted	Draw Above Minimum	Draw At Minimum	Total
Draw Above Minimum	75,886	26,333	102,219
Draw At Minimum	35,701	61,418	97,119
Total	111,587	87,751	199,338

Table 4.6: Draw at Minimum Rate – Binary Choice Classification Diagnostics

Metric	Value
Sensitivity	.700
Specificity	.680
PPV	.632
NPV	.742
Overall Accuracy	.665

Decision to Make Adhoc Drawdown

In this subsection, we use the above procedure to study the binary choice relating to making an adhoc drawdown alongside the income stream generated by one's chosen regular drawdown amounts.

Running the CRE model alongside the PC and RE models for comparison, we can again determine which variables significantly impact on the probability of making an adhoc drawdown. In addition, the signs on the coefficient estimates represent the direction of this influence on the estimated probability relative to increases in the regressor values, although these are not directly interpretable from the raw model output.

As before, in Table 4.8 we make use of `Stata`'s estimation of the average marginal effects to understand broadly how unit changes in the regressor values change the probability of observing an adhoc drawdown.

This time, the age effect is significant in both the linear and quadratic terms. At age 65, an incremental year of age increases the probability of making an adhoc drawdown by approximately 1.2%. The negative sign on the squared age variable creates concavity in the age effect. At the more advanced age of 85, the effect of ageing is lower at 0.7%. The model estimates that at approximately age 115, the marginal effect of ageing would become zero.

Again, although statistically significant, the account balance effect proves to be economically insignificant. The composite effect of the linear and square term becomes positive for non-trivial account balances greater than $200. Even at an account balance of $1,000,000 however, the effect of moving to an account balance of $1,100,000 is only a 0.7% increase in the probability of making an adhoc drawdown.

The minimum drawdown rate has a mild effect, decreasing the probability of an adhoc drawdown by about 0.5% for each increment of 0.01, or 1% of account balance.

Finally, adhoc drawdowns were roughly 4-7% more common until financial year 2008, compared to the base case years of 2014 and 2015.

In-sample classification results, using a cutoff value of 0.5 for predicting a response, are in Table 4.9. The corresponding classification breakdown is in Table 4.10. We see that a cutoff of 0.5 does not provide any sensitivity to true responses. Furthermore, the overall accuracy is not far from what we would expect from using the decision rule 'classify all records as no

Table 4.7: Make Adhoc Drawdown – Binary Choice Regression Model Output

	PC Logit Model	RE Logit Model	CRE Logit Model
Age	0.434***	0.359***	0.438***
	(0.0293)	(0.0638)	(0.102)
Age^2	−0.00445***	−0.00413***	−0.00190**
	(0.000216)	(0.000463)	(0.000671)
Log Account Balance	−0.0670*	0.177**	−0.750***
	(0.0313)	(0.0652)	(0.0865)
(Log Account Balance)^2	−0.00386*	−0.00386	0.0713***
	(0.00157)	(0.00332)	(0.00550)
Minimum Drawdown Rate	−4.011	−7.039*	−7.761*
	(2.114)	(3.279)	(3.903)
Financial Year = 2004	−1.400***	−2.348***	0.685
	(0.0816)	(0.150)	(0.419)
Financial Year = 2005	−1.018***	−1.813***	0.888*
	(0.0790)	(0.143)	(0.382)
Financial Year = 2006	−0.692***	−1.314***	1.080**
	(0.0773)	(0.137)	(0.345)
Financial Year = 2007	−0.490***	−1.035***	1.006**
	(0.0766)	(0.131)	(0.309)
Financial Year = 2008	−0.433***	−0.991***	0.680**
	(0.0496)	(0.0859)	(0.257)
Financial Year = 2009	−1.411***	−2.131***	−0.711**
	(0.0625)	(0.102)	(0.239)
Financial Year = 2010	−1.219***	−1.724***	−0.471*
	(0.0586)	(0.0942)	(0.201)
Financial Year = 2011	−0.546***	−0.992***	−0.0863
	(0.0556)	(0.0886)	(0.167)
Financial Year = 2012	−0.344***	−0.658***	−0.0168
	(0.0361)	(0.0553)	(0.110)
Financial Year = 2013	−0.195***	−0.398***	0.0212
	(0.0365)	(0.0525)	(0.0794)
Risk Appetite	−0.315***	−0.588***	−0.315***
	(0.0338)	(0.0866)	(0.0867)
Gender = Male	0.0721***	0.0947*	0.0993**
	(0.0141)	(0.0370)	(0.0360)
Age at Account Open	0.121***	0.127***	0.161***
	(0.00475)	(0.0108)	(0.0115)
Legacy Account	−0.280***	−0.306**	−1.646***
	(0.0458)	(0.102)	(0.138)
($\bar{\mathbf{x}}_i$ omitted)	.	.	.
Constant	−16.10***	−16.22***	−24.83***
	(0.988)	(2.169)	(2.878)
Observations	205448	205448	204783

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4.8: Make Adhoc Drawdown – Binary Choice Model Average Marginal Effects

	CRE Average Marginal Effects	
Age	0.0274^{***}	(0.00636)
Age^2	-0.000119^{**}	(0.0000420)
Log Account Balance	-0.0470^{***}	(0.00543)
(Log Account Balance)^2	0.00447^{***}	(0.000344)
Minimum Drawdown Rate	-0.486^{*}	(0.245)
Financial Year = 2004	0.0429	(0.0263)
Financial Year = 2005	0.0556^{*}	(0.0239)
Financial Year = 2006	0.0677^{**}	(0.0216)
Financial Year = 2007	0.0630^{**}	(0.0194)
Financial Year = 2008	0.0426^{**}	(0.0161)
Financial Year = 2009	-0.0445^{**}	(0.0149)
Financial Year = 2010	-0.0295^{*}	(0.0126)
Financial Year = 2011	-0.00541	(0.0104)
Financial Year = 2012	-0.00105	(0.00688)
Financial Year = 2013	0.00133	(0.00498)
Risk Appetite	-0.0197^{***}	(0.00543)
Gender = Male	0.00622^{**}	(0.00225)
Age at Account Open	0.0101^{***}	(0.000722)
Legacy Account	-0.103^{***}	(0.00853)
Observations	204783	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

response'—since adhoc drawdown occurs 12.4% of the time.

To try improve the classification results, we tested decision rules based on other cutoff values. However, we did not find a cutoff value that made the classification results more satisfactory.

4.1.3 Continuous Dependent Variable Models

The three (roughly) continuous dependent variables we model are:

1. The excess regular drawdown rate over a financial year, conditional on having drawn above the minimum
2. The unconditional regular drawdown rate over a financial year
3. The adhoc drawdown rate over a financial year, conditional on having made an adhoc drawdown

We fit a sequence of linear panel models where possible: the Pooled Cross-sectional, Fixed and Random Effects, and Hausman-Taylor models. Comparing across these model results sheds insight into the misspecification issues that can be avoided by utilising panel models.

As an inspection of the histograms of these dependent variables and their log transforms will reveal, the first two dependent variables listed are continuous enough to model using linear panel models. By contrast, the third exhibits a significant probability mass, motivating the use of a censored regression model.

Excess Regular Drawdown Rate

We study the excess drawdown rate variable conditional on the retiree drawing above the minimum rates—omitting the probability mass formed at the excess drawdown rate of 0%. Consequently, the results from the models featuring this dependent variable explain the effects of the regressors on the excess drawdown rate, but only for individuals who have elected to draw above the minima.

As a proportion of the account balance, the excess regular drawdown rate is roughly constrained on the interval (0, 1), with the exact upper limit depending on the minimum drawdown rate faced by the individual. For example, an upper limit for the excess regular drawdown rate of 0.95 may be standard for a 65-year-old retiree facing a 5% annual minimum drawdown rate. Since a support of (0, 1) would not be appropriate for a linear model with

Table 4.9: Make Adhoc Drawdown – Binary Choice Classification Table

	Observed		
Predicted	Draw Regular Only	Draw Adhoc	Total
Draw Regular Only	179,599	24,377	203,976
Draw Adhoc	1,271	1,131	2,402
Total	180,870	25,508	206,378

Table 4.10: Make Adhoc Drawdown – Binary Choice Classification Diagnostics

Metric	Value
Sensitivity	.044
Specificity	.993
PPV	.471
NPV	.880
Overall Accuracy	.876

Gaussian errors, we take the natural logarithm of the rates and show the transformed rate histogram in Figure 4.5 and Table 4.11.

The median log excess regular drawdown rate of -2.81 translates to a rate of approximately 6% on the unit scale. We observe a slight peak near 0 on the log scale, corresponding to excess regular drawdown rates nearing 100%. Despite this, we proceed with linear models for this dependent variable, noting that the model may not fit well in the upper tail.

Expanding this pooled histogram through the time dimension in Figures 4.6 and 4.7, we see that the distribution of the transformed dependent variable changes through time.

As these three figures suggest, by incorporating financial year dummy variables to control for financial year effects, it may be reasonable to attempt fitting linear models against this dependent variable.

Table 4.12 provides the model estimation output for the four fitted models for the excess regular drawdown rate. The PC model is *a priori* unlikely to be appropriate, as it assumes there are no unobserved, individual-specific, time-invariant factors that would cause successive observations of the same individual through time to be autocorrelated. By contrast, the FE model removes any time-invariant effects, observed or unobserved, and obtains consistent estimates of the coefficients against the time-varying regressors—the first five regressors and the financial year time dummies. It is in these FE coefficient values that we can be most confident.

On inspecting the RE model coefficients on these time-varying regressors, we notice sizeable discrepancies between the FE and RE models on variables such as Age and Log Account Balance. The Hausman specification test results in Table 4.13 strongly reject the idea that these coefficients are the same at any significance level, and thus we must assume the Random Effects model is unsuitable for drawing inference on the regressors.

The Hausman-Taylor procedure, however, produces estimates of these coefficients which seem much closer to the Fixed Effects model values. Indeed, running the Hausman specification test

Table 4.11: Summary Statistics for Log Excess Regular Drawdown Rate

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Log Excess Regular Drawdown Rate	−5.30	−3.39	−2.81	−2.85	−2.38	0

Figure 4.5: Histogram of Log Excess Regular Drawdown Rate

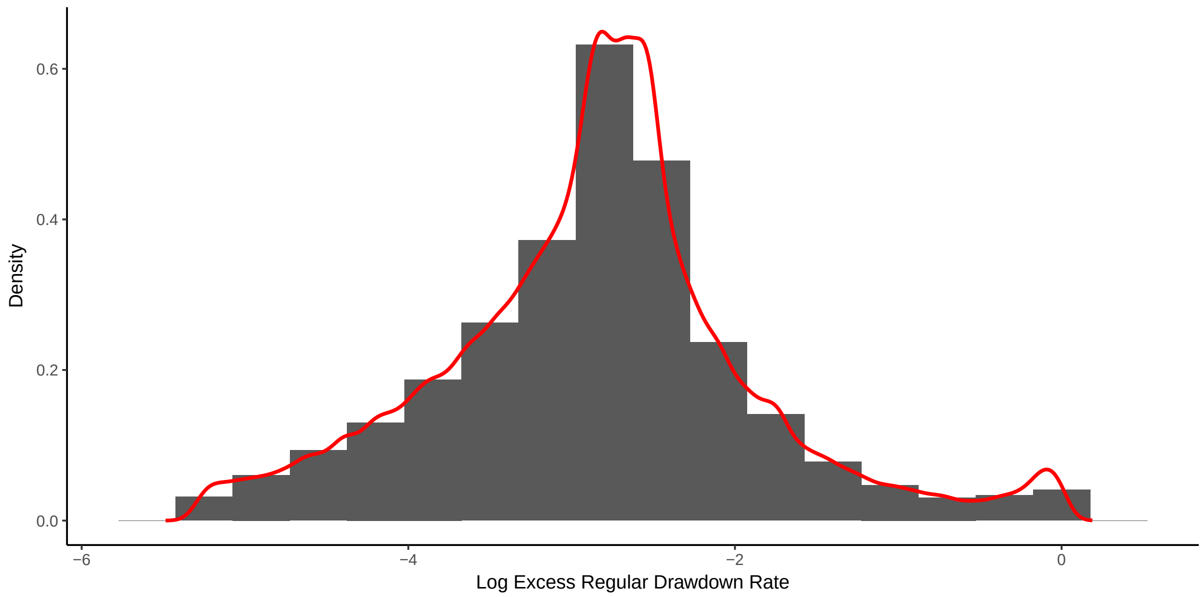

Figure 4.6: 3D Histogram of Log Excess Regular Drawdown Rate over Time

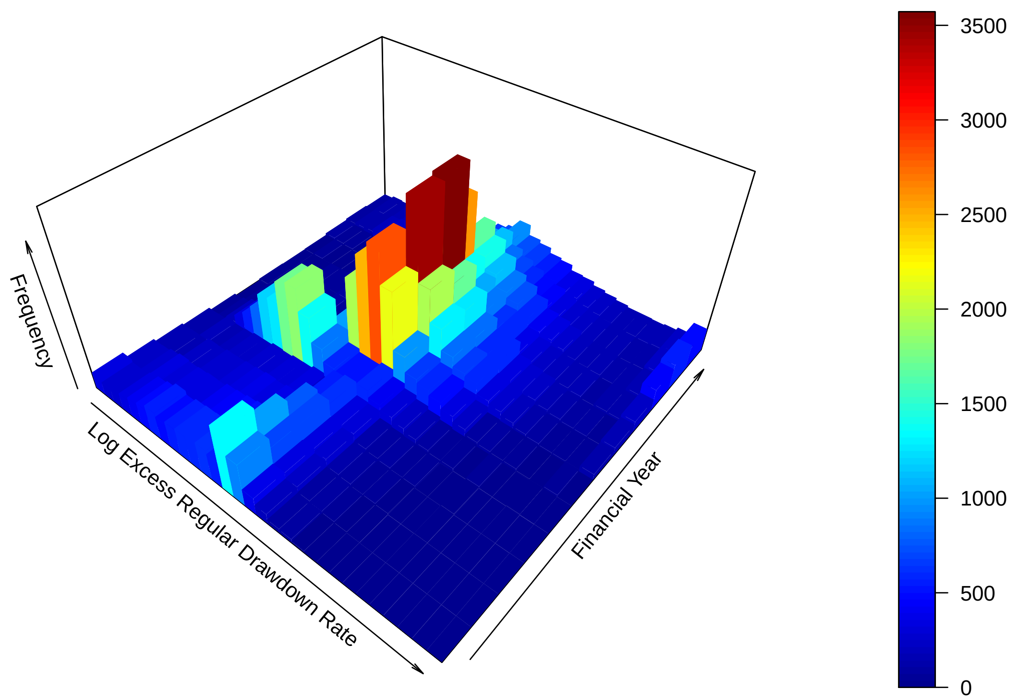

Table 4.12: Log Excess Regular Drawdown Rate – Regression Model Output

	Pooled Cross-Sectional	Fixed Effects	Random Effects	Hausman-Taylor
Age	-0.181^{***}	-0.332^{***}	-0.284^{***}	-0.326^{***}
	(0.00649)	(0.0177)	(0.0131)	(0.0177)
Age^2	0.00141^{***}	0.00221^{***}	0.00205^{***}	0.00220^{***}
	(0.0000466)	(0.000114)	(0.0000934)	(0.000114)
Log Account Balance	-0.493^{***}	0.741^{***}	0.319^{***}	0.744^{***}
	(0.0189)	(0.0296)	(0.0194)	(0.0297)
$(\text{Log Account Balance})^2$	0.00579^{***}	-0.0671^{***}	-0.0388^{***}	-0.0674^{***}
	(0.000891)	(0.00187)	(0.00108)	(0.00187)
Minimum Drawdown Rate	-1.384^{*}	-3.510^{***}	-3.904^{***}	-3.513^{***}
	(0.610)	(0.591)	(0.579)	(0.591)
Financial Year = 2004	-0.557^{***}	-0.982^{***}	-0.855^{***}	-0.938^{***}
	(0.0202)	(0.0626)	(0.0273)	(0.0615)
Financial Year = 2005	-0.534^{***}	-1.007^{***}	-0.895^{***}	-0.967^{***}
	(0.0205)	(0.0572)	(0.0260)	(0.0563)
Financial Year = 2006	-0.449^{***}	-0.963^{***}	-0.873^{***}	-0.927^{***}
	(0.0207)	(0.0521)	(0.0248)	(0.0513)
Financial Year = 2007	-0.409^{***}	-0.935^{***}	-0.868^{***}	-0.902^{***}
	(0.0214)	(0.0471)	(0.0239)	(0.0464)
Financial Year = 2008	-0.579^{***}	-0.564^{***}	-0.585^{***}	-0.537^{***}
	(0.0116)	(0.0383)	(0.0160)	(0.0377)
Financial Year = 2009	0.144^{***}	0.0662	0.0486^{*}	0.0887^{*}
	(0.0189)	(0.0354)	(0.0208)	(0.0349)
Financial Year = 2010	0.349^{***}	0.169^{***}	0.194^{***}	0.187^{***}
	(0.0187)	(0.0300)	(0.0198)	(0.0296)
Financial Year = 2011	0.343^{***}	0.195^{***}	0.206^{***}	0.210^{***}
	(0.0187)	(0.0254)	(0.0188)	(0.0252)
Financial Year = 2012	0.207^{***}	0.0940^{***}	0.104^{***}	0.105^{***}
	(0.0128)	(0.0164)	(0.0115)	(0.0162)
Financial Year = 2013	0.245^{***}	0.163^{***}	0.176^{***}	0.170^{***}
	(0.0133)	(0.0119)	(0.0105)	(0.0118)
Risk Appetite	0.117^{***}		0.0981^{***}	0.169^{**}
	(0.0108)		(0.0231)	(0.0583)
Gender = Male	0.113^{***}		0.145^{***}	0.449^{***}
	(0.00436)		(0.00958)	(0.0270)
Age at Account Open	-0.00734^{***}		0.00134	-0.297^{***}
	(0.00111)		(0.00238)	(0.0159)
Legacy Account	-0.0428^{***}		-0.0247	-0.312^{***}
	(0.0112)		(0.0232)	(0.0619)
Constant	7.995^{***}	9.822^{***}	8.188^{***}	28.13^{***}
	(0.237)	(0.752)	(0.435)	(1.085)
Observations	111585	111585	111585	111585

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4.13: Log Excess Regular Drawdown Rate – Hausman Test: FE vs RE

Metric	Value
χ^2_{14} Test Statistic	8198.48
p-value	0.0000

Figure 4.7: Heatmap of Log Excess Regular Drawdown Rate over Time

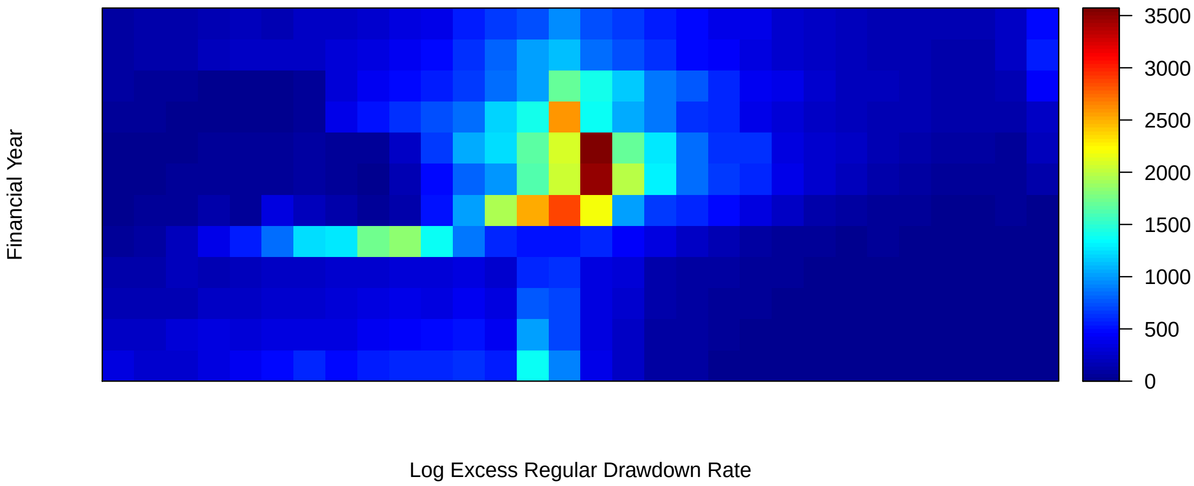

between these two models fails to reject the null hypothesis that these coefficients differ systematically, at any significance level smaller than 91%. Results from this second Hausman test are in Table 4.14.

This failure to reject implies that the HT model may not be inappropriate, and so we use these coefficient estimates to perform inference on all regressors, including the time-invariant $\mathbf{z}_i$.

Positive signs represent effects that increase the conditional expectation of the dependent variable as the corresponding regressor increases, and vice-versa. As the dependent variable is the log-transformed excess regular drawdown rate, the values of the coefficients reflect proportional changes in the excess regular drawdown rate for unit increases in the regressor values. For example, comparing a male and a female who are identical in all other regressors and who both make drawdowns above the minimum, the excess drawdown rate for the male can be expected to be that of his female counterpart multiplied by approximately $e^{0.449} = 1.57$. Similarly, for each year of age retirement is delayed, drawdown rates in excess of the minimum tend to reduce through multiplication by a factor of approximately $e^{-0.297} = 0.74$. While these proportional factors appear large in magnitude, they apply only to the rate of excess drawdown, which has a median value of 6%.

An incremental year of age for a 65-year-old scales the expected excess regular drawdown rate

Table 4.14: Log Excess Regular Drawdown Rate – Hausman Test: FE vs HT

Metric	Value
χ^2_{14} Test Statistic	7.54
p-value	0.9119

down by 4%, while the same increase at age 85 scales the expected rate up by approximately 5%. The turning point for this parabolic effect occurs around age 74. Comparatively, the effect of account balance is to decrease the excess regular drawdown rate, as the negative coefficient on the square term dominates for any account balances greater than $250. At an account balance of $100,000, an increase in the account balance by 10% to $110,000 scales the expected drawdown rate down by approximately 7.7%. At a balance of $1,000,000, the same proportional increase results in a scaling down by 10.7%.

Other notable results include: the negative signs of financial years 2004–7, where drawdown at the minimum was more common; the negative sign of financial year 2008, where there was a spike in the number of small excess drawdowns—corresponding to individuals leaving the minima for the first time; the negative sign for legacy accounts, which are more likely to draw at the minima; and the positive effect of increasing risk appetite.

Aside from inferring the effects of the available regressors on excess regular drawdown rates, we also inspect some model diagnostic plots in Figure 4.8, to determine how well the model assumptions are satisfied.

The top two panels indicate that the residuals are not exactly Normally distributed, which means that the model assumptions do not hold precisely. In addition, the plot of residuals vs. fitted values shows a prominent linear trend, whereas the ideal plot would have a horizontal trend. Plotting residuals against individual explanatory variables shows multiple instances of heteroscedasticity and linear trending. Overall, we are convinced that there are still relevant, omitted variables that are correlated with at least some of our regressors and have a significant effect on the dependent variable.

Regular Drawdown Rate

Our second set of continuous dependent variable models examine the effect of the available regressors on the unconditional regular drawdown rate—that is, including all individuals drawing at and above the minimum drawdown rates.

The regular drawdown rate is bounded from above by 100%, corresponding to a value of zero on the log scale. Additionally, the smallest concessional minimum drawdown rate attainable during the sample period was 2%, for a retiree younger than 65 and during financial years 2009–11, inclusive. Furthermore, where a retiree makes one or more adhoc drawdowns during a financial year, they can reduce their regular drawdown amounts such that the annual regular drawdown rate is less than the legislated minimum rates, while still keeping their total rate of drawdown at or above the minima. Total drawdown rates below the minima are possible, but likely to be rare because they attract penalties through the taxation system. Figure 4.9 and Table 4.15 reveal the distribution of the log regular drawdown rates.

Here, the median value of -2.57 on the log scale corresponds to drawdown rates of approximately 7.7% on the unit scale. Inspection suggests that the probability masses at the endpoints are mild. We proceed with linear modelling techniques, although aware that the fit in the tails

Figure 4.8: Log Excess Regular Drawdown Rate – Hausman-Taylor Model Residual Diagnostics

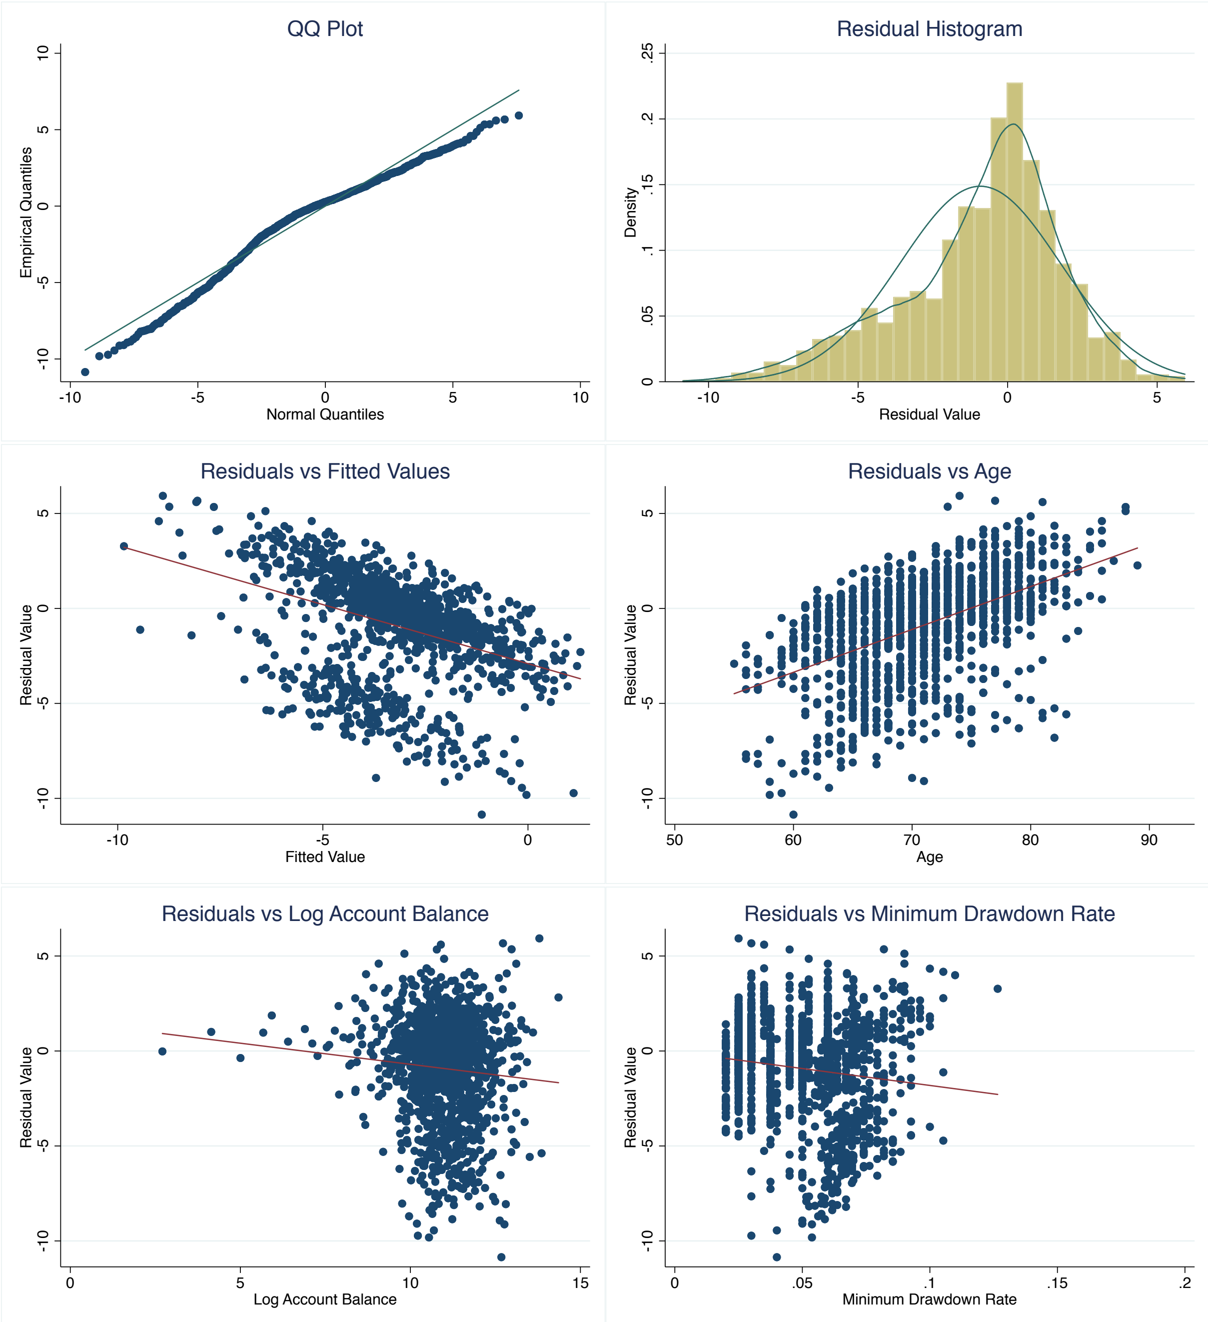

Table 4.15: Summary Statistics for Log Regular Drawdown Rate

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Log Regular Drawdown Rate	-9.37	-2.82	-2.57	-2.50	-2.23	0

Figure 4.9: Histogram of Log Regular Drawdown Rate

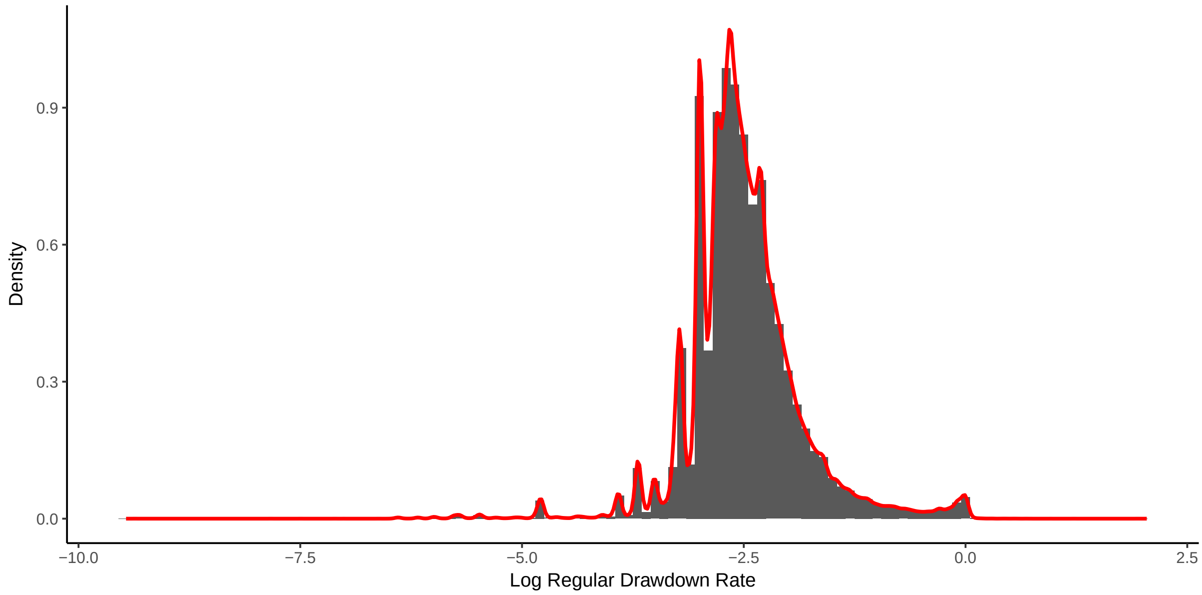

may suffer as a result.

Figures 4.10 and 4.11 spread the dependent variable over the time index. In this case, the location of the dependent variable seems stable, although its spread appears to increase through time.

Table 4.16 contains the regression estimates of the PC, FE, RE and HT models for the log regular drawdown rate. Again, we assume the FE estimates to be the most correct, and include the PC model output for completeness only.

As before, the Hausman test result provided in Table 4.17 strongly rejects, at any level of significance, the hypothesis that the FE and RE model estimates are the same. By contrast, Table 4.18 shows that the Hausman test fails to reject the same hypothesis when comparing the FE and HT estimates, at any level of significance. Consequently, we use the HT model results for inference on the effects of the regressors on the regular drawdown rate.

Except for the 2013 financial year effect, all regressors are statistically significant at least at the 5% level. After controlling for other available regressors, drawdown rates are, on average, higher for retirees with higher derived risk appetites, and for individuals retiring older. Compared to their female counterparts, male retirees draw at rates that are larger by a multiplicative factor of 1.13. Naturally, increasing minimum drawdown rates cause expected drawdown rates to increase.

The effect of increasing age over time is broadly negative, but with some convexity. An incremental year of age for a 65-year-old decreases the expected regular drawdown rate by a scaling factor of about 2.9%, while the same increment at age 85 only decreases the expected drawdown rate by a scaling factor of 0.1%. At age 86, the effect is zero.

Despite the positive coefficient on the linear age term, the quadratic effect dominates for all

Figure 4.10: 3D Histogram of Log Regular Drawdown Rate over Time

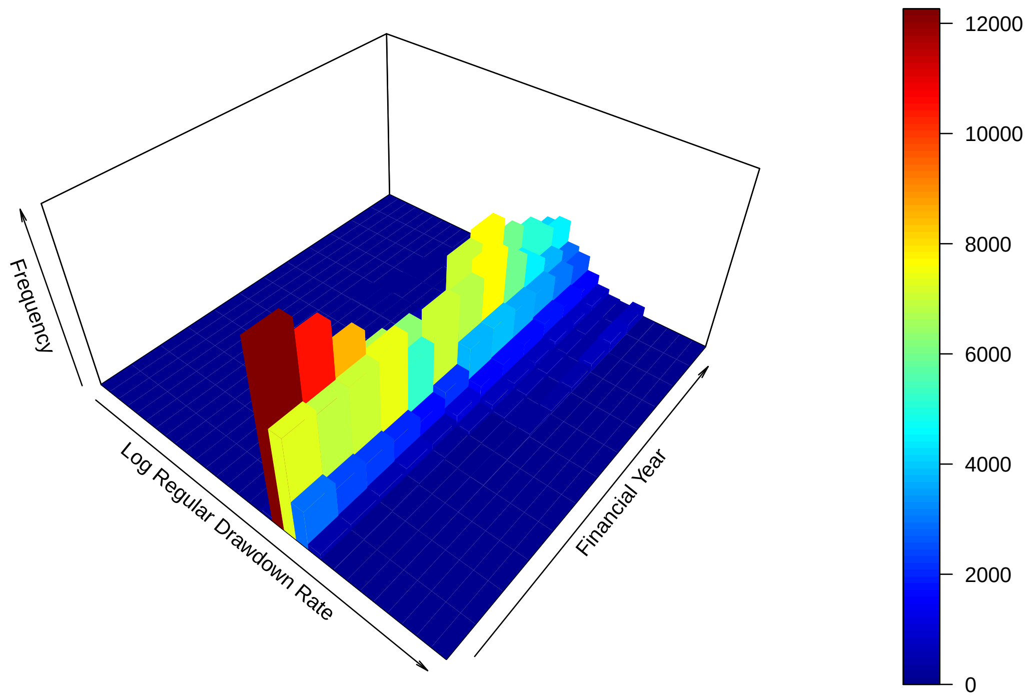

Figure 4.11: Heatmap of Log Regular Drawdown Rate over Time

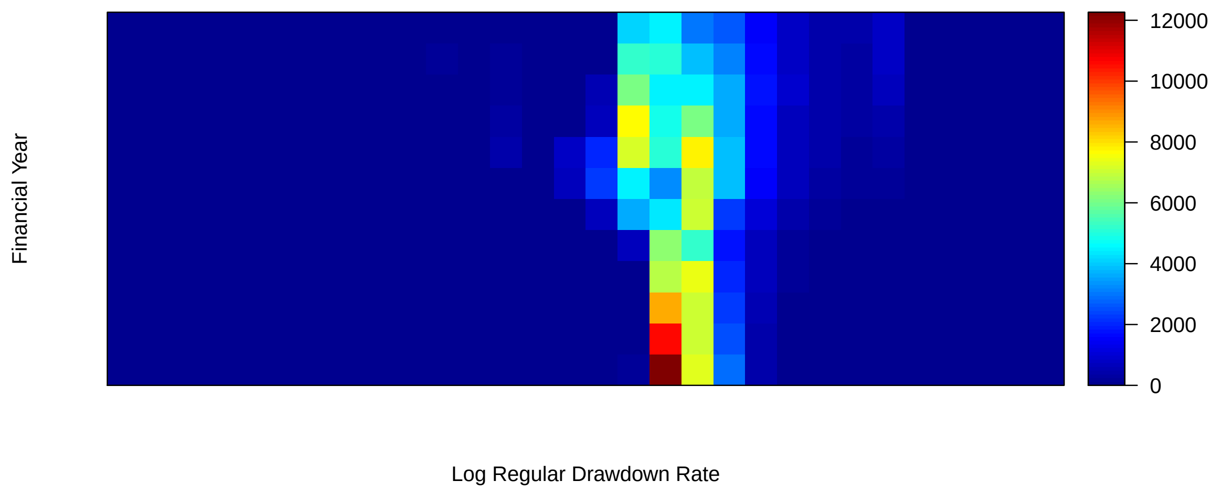

Table 4.16: Log Regular Drawdown Rate – Regression Model Output

	Pooled Cross-Sectional	Fixed Effects	Random Effects	Hausman-Taylor
Age	-0.0425^{***}	-0.118^{***}	-0.0815^{***}	-0.118^{***}
	(0.00390)	(0.00869)	(0.00703)	(0.00869)
Age^2	0.000428^{***}	0.000684^{***}	0.000621^{***}	0.000684^{***}
	(0.0000282)	(0.0000553)	(0.0000519)	(0.0000553)
Log Account Balance	-0.673^{***}	0.220^{***}	-0.0507^{**}	0.220^{***}
	(0.0160)	(0.0229)	(0.0174)	(0.0229)
$(\text{Log Account Balance})^2$	0.0218^{***}	-0.0299^{***}	-0.0123^{***}	-0.0299^{***}
	(0.000727)	(0.00140)	(0.000955)	(0.00140)
Minimum Drawdown Rate	4.805^{***}	3.524^{***}	3.323^{***}	3.524^{***}
	(0.251)	(0.310)	(0.301)	(0.310)
Financial Year = 2004	-0.136^{***}	-0.459^{***}	-0.231^{***}	-0.459^{***}
	(0.0103)	(0.0344)	(0.0165)	(0.0344)
Financial Year = 2005	-0.150^{***}	-0.427^{***}	-0.227^{***}	-0.427^{***}
	(0.0101)	(0.0315)	(0.0157)	(0.0315)
Financial Year = 2006	-0.146^{***}	-0.376^{***}	-0.206^{***}	-0.376^{***}
	(0.0100)	(0.0286)	(0.0150)	(0.0286)
Financial Year = 2007	-0.146^{***}	-0.326^{***}	-0.185^{***}	-0.326^{***}
	(0.00988)	(0.0259)	(0.0143)	(0.0259)
Financial Year = 2008	-0.0804^{***}	-0.242^{***}	-0.138^{***}	-0.242^{***}
	(0.00616)	(0.0214)	(0.00904)	(0.0214)
Financial Year = 2009	0.0950^{***}	-0.120^{***}	-0.0262^{**}	-0.120^{***}
	(0.00783)	(0.0191)	(0.00913)	(0.0191)
Financial Year = 2010	0.103^{***}	-0.126^{***}	-0.0310^{***}	-0.126^{***}
	(0.00818)	(0.0159)	(0.00891)	(0.0159)
Financial Year = 2011	0.0655^{***}	-0.0806^{***}	-0.0156	-0.0806^{***}
	(0.00833)	(0.0134)	(0.00847)	(0.0134)
Financial Year = 2012	0.0412^{***}	-0.0541^{***}	-0.00806	-0.0541^{***}
	(0.00644)	(0.00868)	(0.00518)	(0.00868)
Financial Year = 2013	0.0659^{***}	-0.0116	0.0225^{***}	-0.0116
	(0.00690)	(0.00610)	(0.00478)	(0.00610)
Risk Appetite	0.0715^{***}	0	0.0876^{***}	0.0996^{***}
	(0.00626)	(.)	(0.0150)	(0.0158)
Gender = Male	0.0799^{***}	0	0.109^{***}	0.118^{***}
	(0.00238)	(.)	(0.00592)	(0.00623)
Age at Account Open	-0.00732^{***}	0	0.00394^{**}	0.0306^{***}
	(0.000555)	(.)	(0.00128)	(0.00395)
Legacy Account	0.181^{***}	0	0.188^{***}	0.297^{***}
	(0.00610)	(.)	(0.0130)	(0.0342)
Constant	3.184^{***}	3.693^{***}	1.696^{***}	1.437^{***}
	(0.155)	(0.383)	(0.238)	(0.261)
Observations	204221	204221	204221	204221

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4.17: Log Regular Drawdown Rate – Hausman Test: FE vs RE

Metric	Value
χ^2_{14} Test Statistic	6499.40
p-value	0.0000

Table 4.18: Log Regular Drawdown Rate – Hausman Test: FE vs HT

Metric	Value
χ^2_{14} Test Statistic	0.00
p-value	1.0000

nontrivial account balances. At a balance of \$100,000, a 10% increase in account balance drives a reduction in the drawdown rate by a factor of 4.5%, while the same proportional increment at a balance of \$1,000,000 scales the drawdown rate down by a factor of 5.8%.

In general, retirees drew at higher rates in later financial years, with financial years before 2008 exhibiting substantially lower relative rates.

We turn to the residual diagnostics in Figure 4.12. The empirical distribution of residuals shows a left tail much heavier than a comparison Normal distribution. Moreover, the residual series plotted against fitted values and some regressors indicate there are still unobserved, relevant factors that our model is not incorporating. Specifically, at smaller account balances, we systematically overestimate the drawdown rate, and vice versa for higher account balances.

Adhoc Drawdown Rate

Finally, we examine the rate at which adhoc drawdowns deplete account balances, for those who use their account to make adhoc withdrawals. Figure 4.13 and Table 4.19 describe the distribution of the adhoc drawdown rate after taking the natural logarithm.

The median value of -2.27 on the log scale translates to a drawdown rate of approximately 10%. However, the most interesting feature of this distribution is the significant probability mass sitting at a value close to 0 on the log scale, corresponding to a complete withdrawal of the account balance as a lump sum. Roughly 8% of all adhoc drawdowns are used to completely withdraw the account balance out of the superannuation system. This probability mass near zero motivates our subsequent use of a censored regression model.

Spreading the log adhoc drawdown rate through the time dimension in Figures 4.14 and 4.15, we observe that over time, an increasing number of adhoc drawdowns draw down the entire account balance.

We use the tobit censored regression model, estimating by CRE. For comparison purposes, Table 4.20 provides the PC and standard RE model coefficients alongside the CRE estimates.

These coefficients are the marginal effects on the dependent variable relative to unit increases in the corresponding regressors. As the dependent variable is on the log scale, these coeffi-

Table 4.19: Summary Statistics for Log Adhoc Drawdown Rate

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Log Adhoc Drawdown Rate	−12.88	−3.15	−2.27	−2.28	−1.15	−0.11

Figure 4.12: Log Regular Drawdown Rate – Hausman-Taylor Model Residual Diagnostics

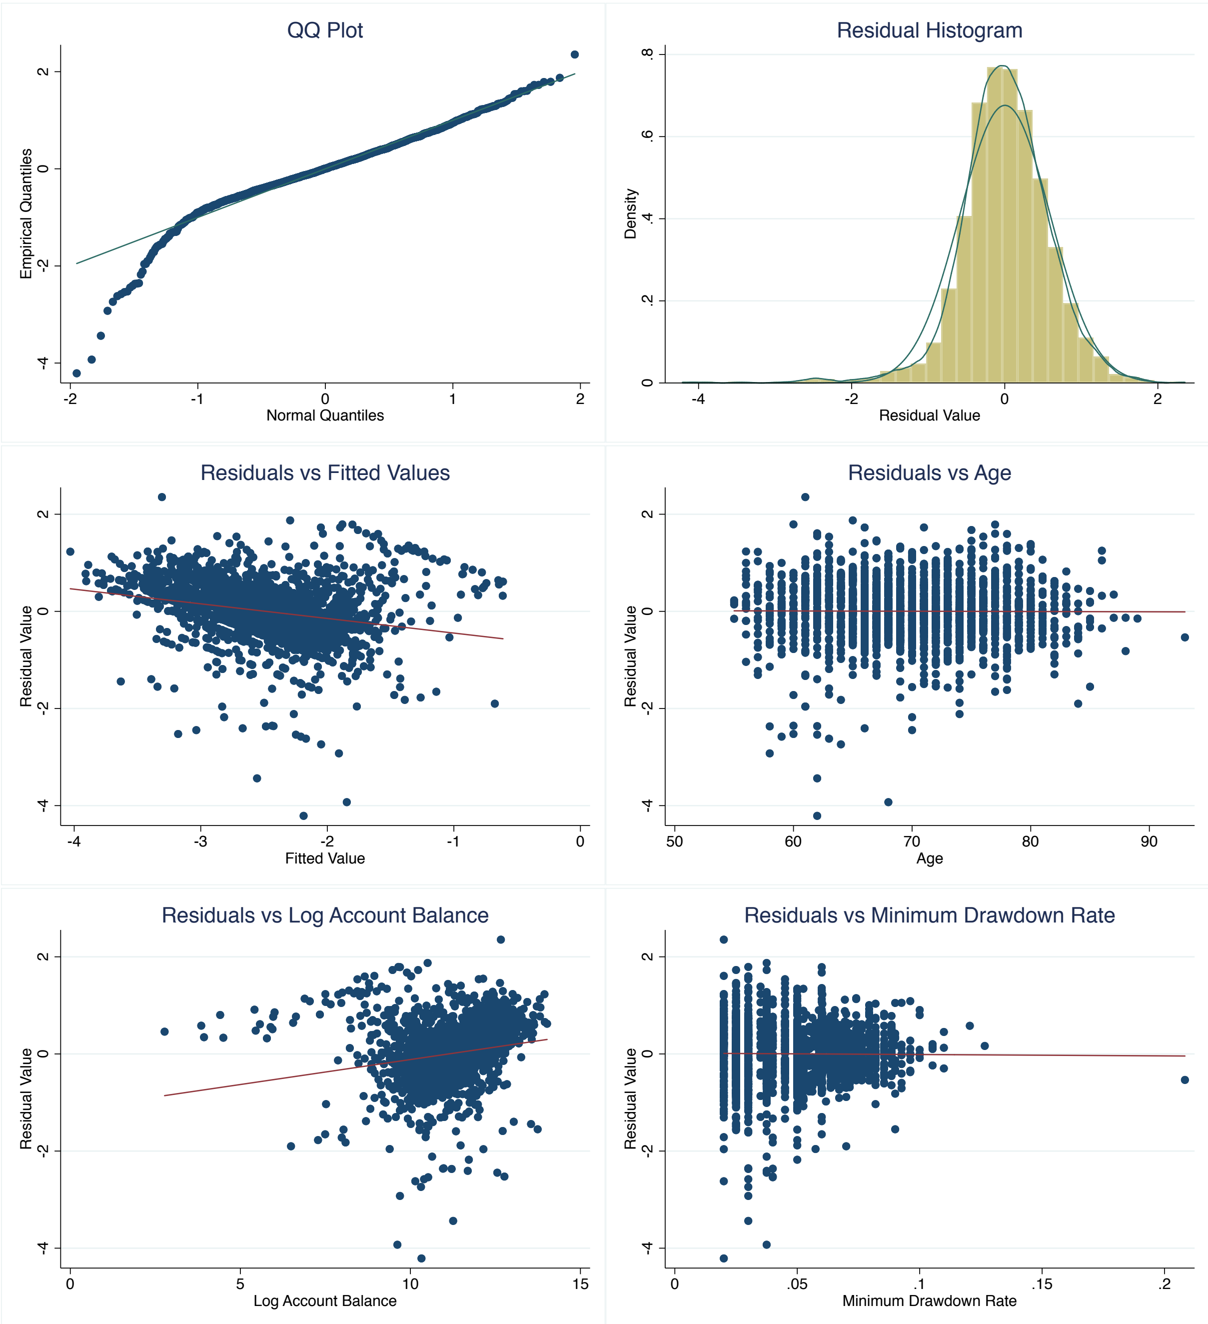

Figure 4.13: Histogram of Log Adhoc Drawdown Rate

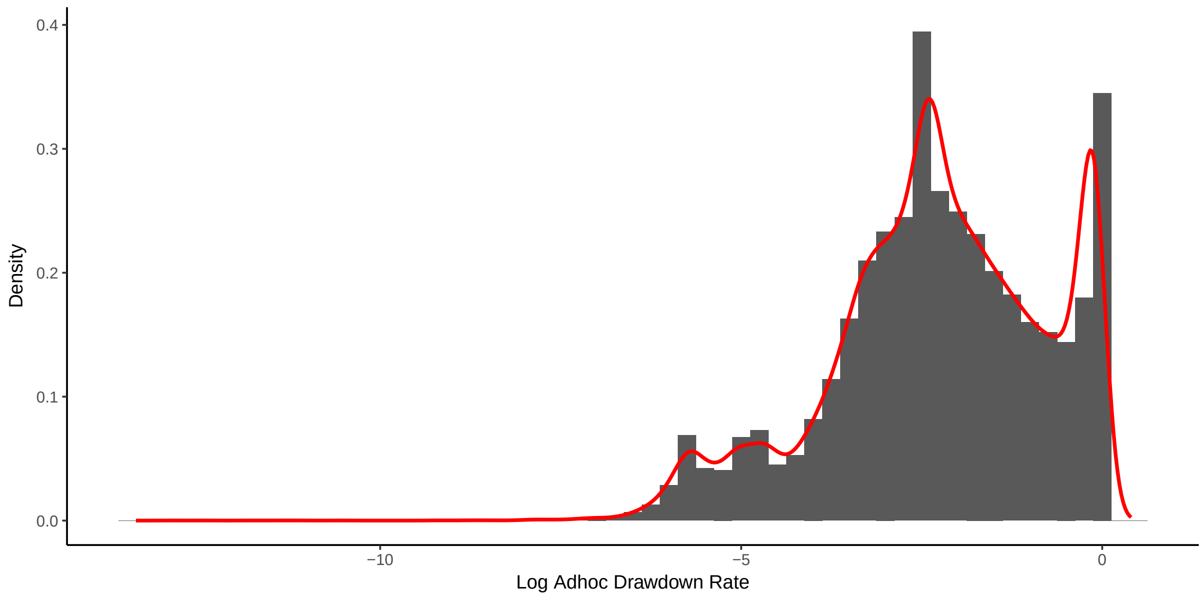

Figure 4.14: 3D Histogram of Log Adhoc Drawdown Rate over Time

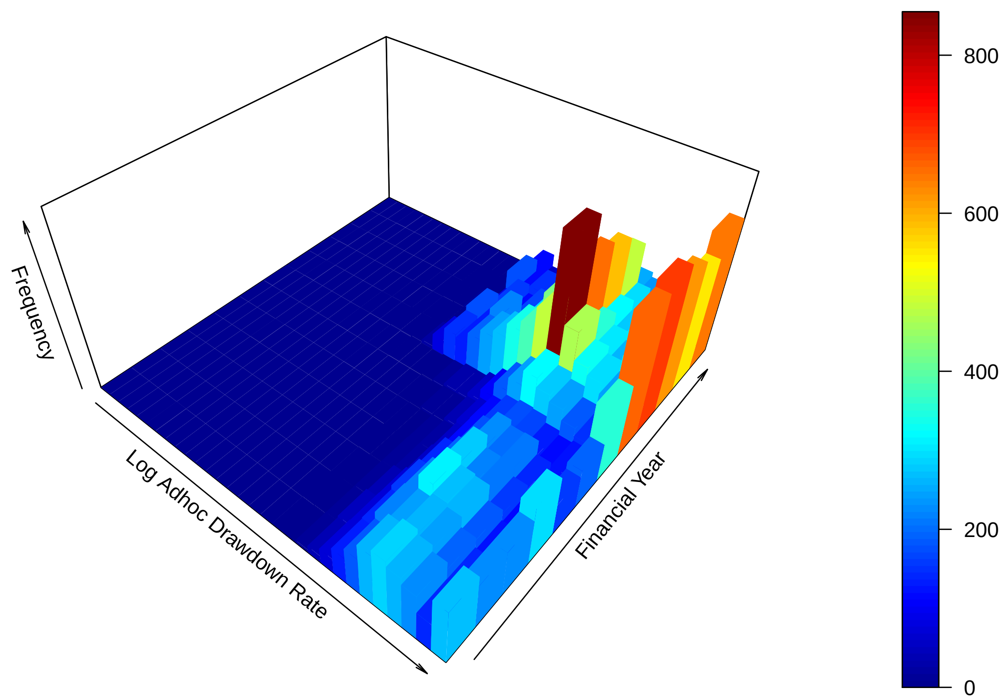

Table 4.20: Log Adhoc Drawdown Rate – Regression Model Output

	PC Tobit Model	RE Tobit Model	CRE Tobit Model
Age	-0.0909^{**}	-0.0198	0.125^{*}
	(0.0312)	(0.0349)	(0.0550)
Age^2	0.00128^{***}	0.000817^{**}	-0.000281
	(0.000232)	(0.000252)	(0.000339)
Log Account Balance	0.189^{***}	0.212^{***}	-0.275^{***}
	(0.0335)	(0.0299)	(0.0342)
(Log Account Balance)^2	-0.0223^{***}	-0.0115^{***}	0.0287^{***}
	(0.00168)	(0.00162)	(0.00200)
Minimum Drawdown Rate	-18.81^{***}	-7.934^{***}	-6.215^{**}
	(2.649)	(2.231)	(2.393)
Financial Year = 2004	0.865^{***}	0.144	-0.842^{**}
	(0.101)	(0.104)	(0.293)
Financial Year = 2005	1.018^{***}	0.238^{*}	-0.708^{**}
	(0.0977)	(0.0984)	(0.267)
Financial Year = 2006	0.975^{***}	0.235^{*}	-0.643^{**}
	(0.0964)	(0.0942)	(0.240)
Financial Year = 2007	1.042^{***}	0.308^{***}	-0.519^{*}
	(0.0963)	(0.0907)	(0.215)
Financial Year = 2008	0.464^{***}	0.0620	-0.665^{***}
	(0.0654)	(0.0629)	(0.180)
Financial Year = 2009	0.116	-0.198^{**}	-0.887^{***}
	(0.0808)	(0.0715)	(0.162)
Financial Year = 2010	0.236^{**}	-0.0275	-0.485^{***}
	(0.0744)	(0.0648)	(0.135)
Financial Year = 2011	-0.298^{***}	-0.152^{**}	-0.399^{***}
	(0.0687)	(0.0587)	(0.109)
Financial Year = 2012	0.0188	-0.0121	-0.210^{**}
	(0.0451)	(0.0366)	(0.0734)
Financial Year = 2013	-0.0646	0.00369	-0.0630
	(0.0455)	(0.0348)	(0.0514)
Risk Appetite	-0.273^{***}	-0.423^{***}	-0.111^{*}
	(0.0371)	(0.0527)	(0.0482)
Gender = Male	0.183^{***}	0.101^{***}	0.111^{***}
	(0.0182)	(0.0259)	(0.0229)
Age at Account Open	-0.0567^{***}	-0.0772^{***}	0.0176^{*}
	(0.00592)	(0.00790)	(0.00766)
Legacy Account	-0.181^{**}	0.155^{*}	-0.343^{***}
	(0.0580)	(0.0745)	(0.0930)
($\bar{\mathbf{x}}_i$ omitted)	.	.	.
Constant	3.039^{**}	-0.0839	8.027^{***}
	(1.033)	(1.141)	(1.600)
σ_α		1.210^{***}	1.026^{***}
		(0.0107)	(0.00963)
σ_e		0.802^{***}	0.792^{***}
		(0.00504)	(0.00490)
Observations	25076	25076	24947

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 4.15: Heatmap of Log Adhoc Drawdown Rate over Time

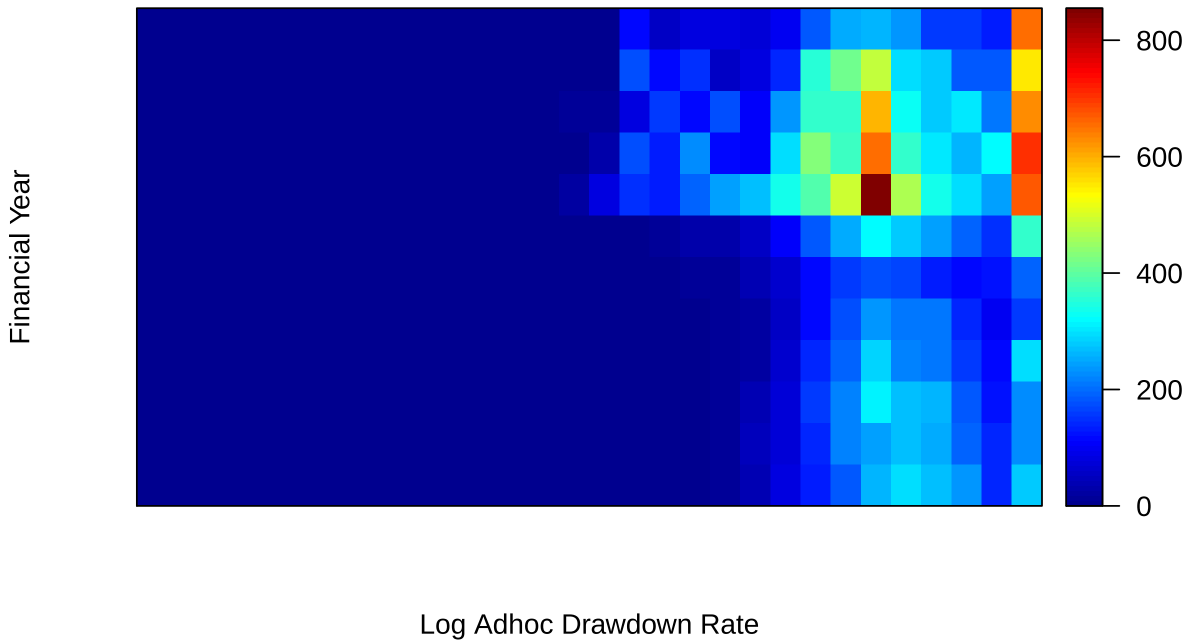

cients represent proportional changes in the adhoc drawdown rate. For example, increasing the minimum drawdown rate by 0.01, or 1% of account balance, multiplies the expected drawdown rate by a factor of approximately 0.94.

The effect of ageing is only significant in the linear term. An incremental year of age scales the expected adhoc drawdown rate up by a factor of 12.5%. The square term on the log account balance variable dominates for all reasonable account balances, in the positive direction. At \$100,000, a 10% increase in account balance drives an expected 3.6% proportionally larger drawdown rate, while at a balance of \$1,000,000 a 10% increase can expect to scale the adhoc drawdown rate up by 4.9%.

In Figure 4.16, we inspect residual diagnostic plots for the current model. There is heteroscedasticity in the residuals with respect to the log account balance, but broadly these diagnostics seem better than in the previous two continuous dependent variable models.

4.1.4 Summary of Panel Regression Model Results

Overall, the panel regression models estimate statistical relationships between several dependent variables of interest and the available regressors. These include the binary choice observations of drawing at the minimum rates and making an adhoc drawdown, as well as models for the rate of regular and adhoc drawdowns.

Figure 4.16: Log Adhoc Drawdown Rate – CRE Tobit Model Residual Diagnostics

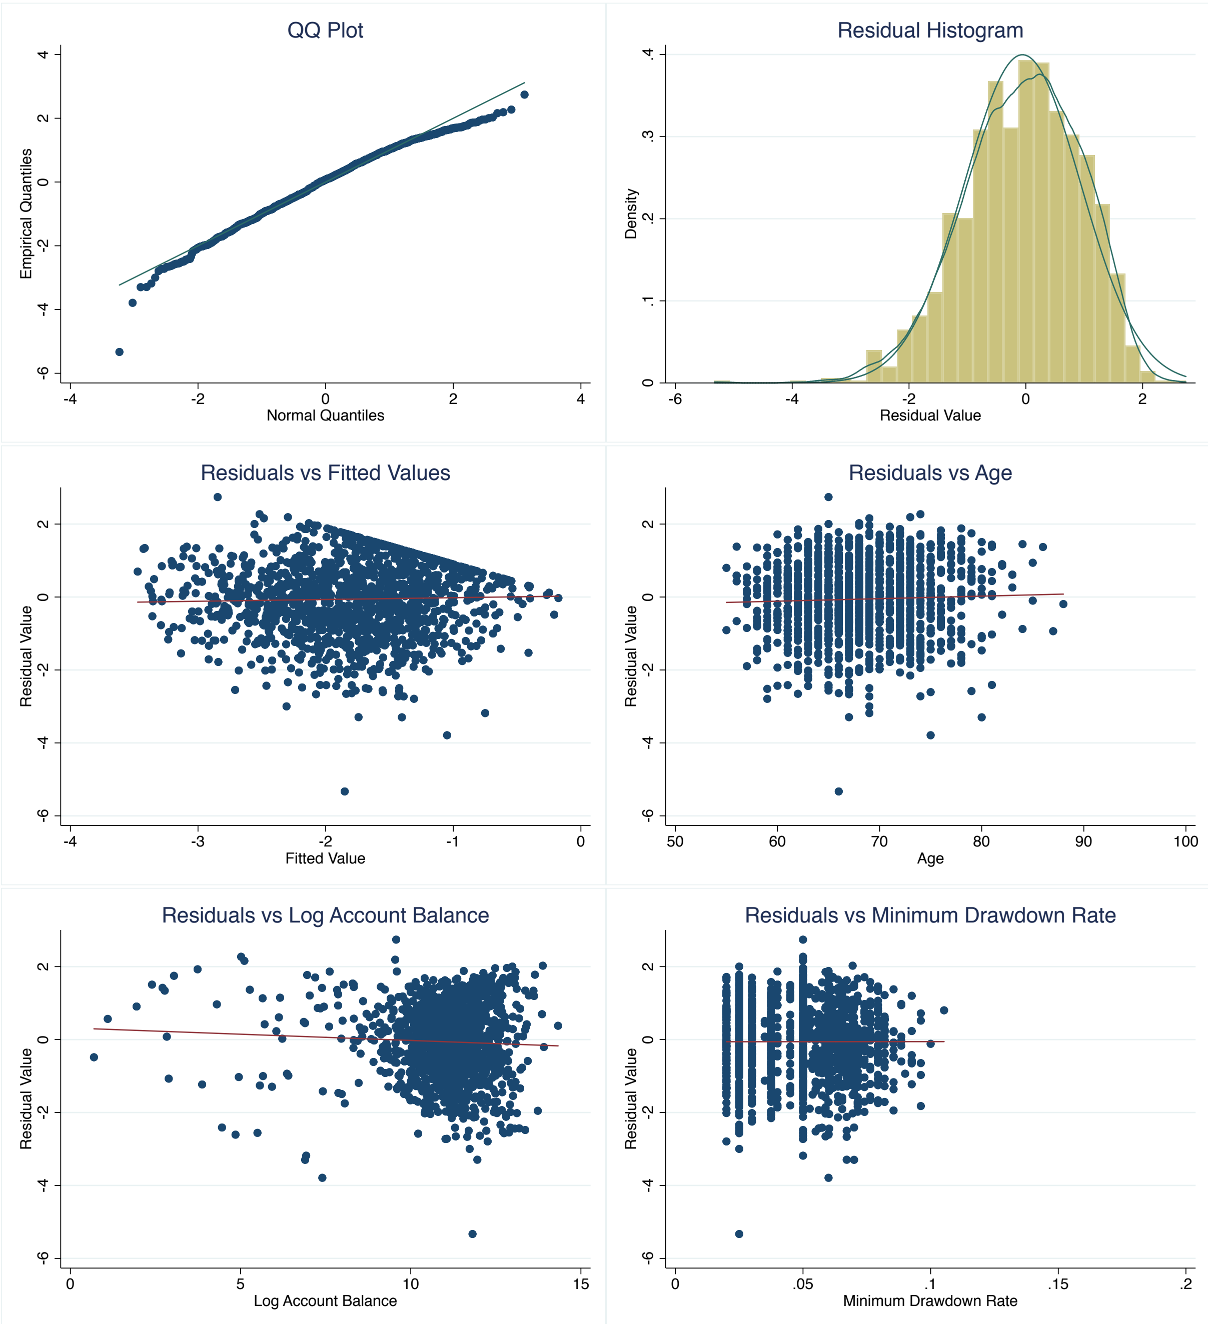

Although these models provide insights into the impact of the regressors on drawdowns, the model diagnostics convince us that much of the variation in observed outcomes remains uncaptured using administrative data alone. Motivated by hypotheses drawn from the theoretical literature, and permitted by the panel nature of our data, we proceed to study drawdowns over time by adding a behavioural dimension to the analysis.

4.2 Component 2: Cluster Analysis

In this section, panel data visualisations inspire the manual and machine-assisted procedures to identify groups that are similar in their observed drawdown behaviours over time.

4.2.1 Panel Visualisations

A quantity of particular interest is the rate at which individuals intend to draw down their accounts, exclusive of their adhoc drawdowns. Figure 4.17 shows the regular drawdown rates for all 44,000 accounts which joined the sample in financial years 2004 and 2009–11. Each line segment represents an individual's trajectory in the dependent variable through time.

Immediately, two aspects to the data become clear. First, since the latest minimum drawdown rules came into effect on 1 July 2007, many individuals are able to draw from their accounts at constant rates. Second, groups suggest themselves visually, through inspection alone.

As individuals face different minimum drawdown requirements at different ages and in different financial years, Figure 4.18 visualises how these rates translate into excess regular drawdown rates. Here, people at the zero line are drawing exactly at their respective minimum rates, while drawdowns above the minimum have a nonzero value on the vertical axis.

In making decisions as to their regular drawdowns, individuals may focus on the dollar amount taken, rather than the rate this represents. The regular drawdown amount for all accounts, in nominal dollar terms, is given in Figure 4.19. This plot makes visible the tendency for many retirees to draw down level amounts over time.

In determining a drawdown rate, the account balance as the denominator is highly influential—especially for individuals who tend towards level drawdown amounts over time. To observe any patterns in account balances over time that would directly affect drawdown rates, Figure 4.20 is useful. The decline in the balances for 2009 and 2010 is evident. Note, however, that as the account balances plotted are as at the start of the relevant financial year, the corresponding declines in account balances occurred during financial years ended 30 June 2008 and 2009. This effect may have at least partly caused observed increases in the regular—and excess regular—drawdown rates plotted for financial years 2009 and 2010. In the prior regression modelling, including the account balance and financial year dummies as regressors controlled for the influence of the account balances and financial year-specific effects on the drawdown rates. This allowed the inference on the remaining regressors to be free from these effects.

Figure 4.17: Regular Drawdown Rate Panel Visualisation

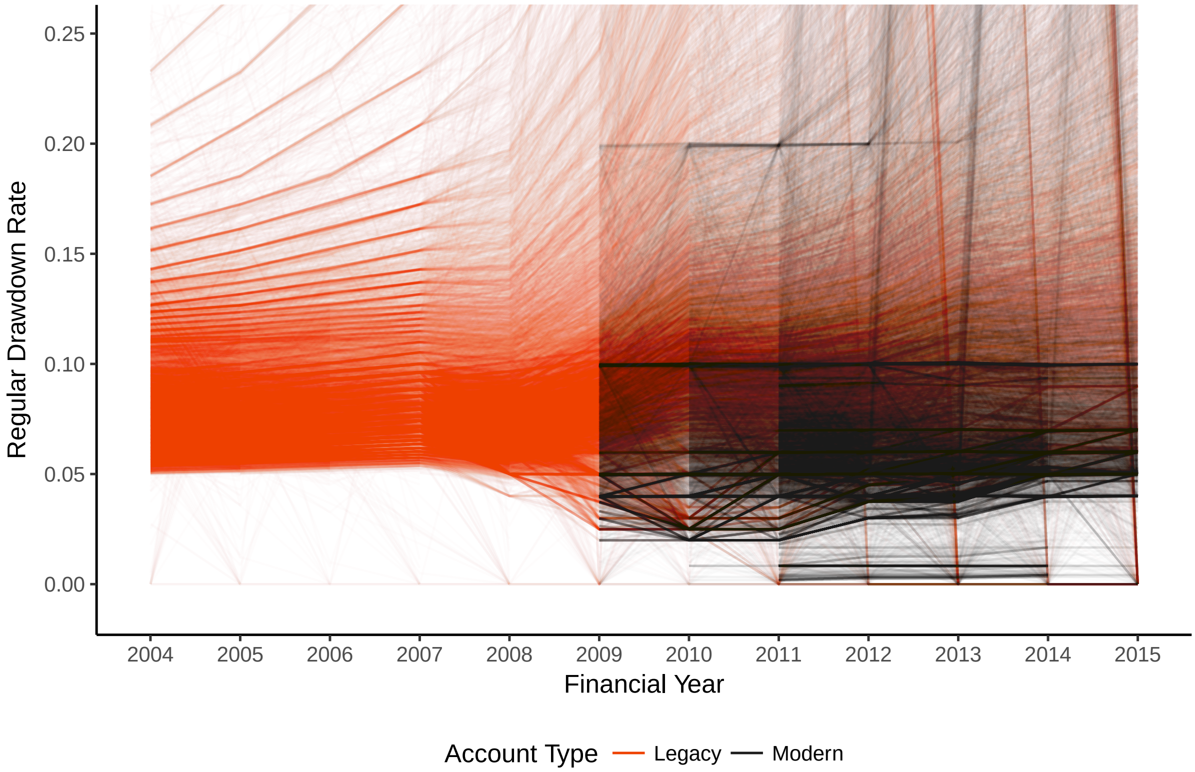

Figure 4.18: Excess Drawdown Rate Panel Visualisation

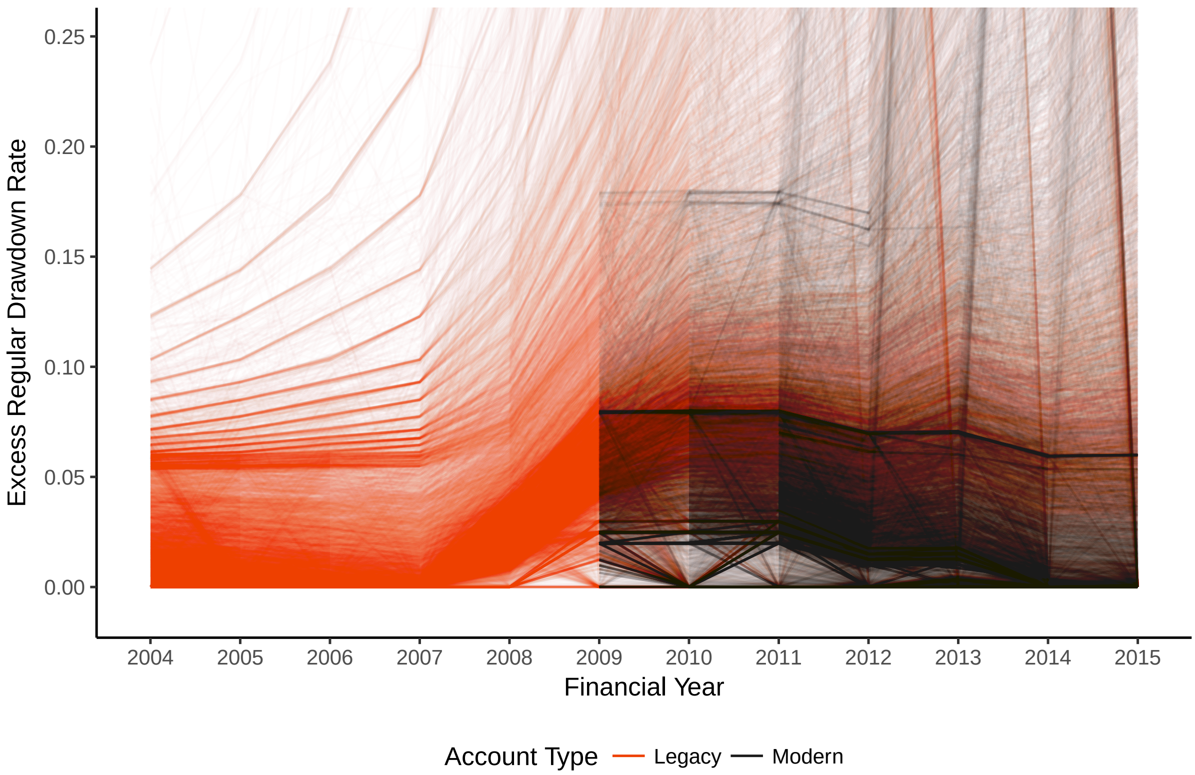

Figure 4.19: Regular Drawdown Amount Panel Visualisation

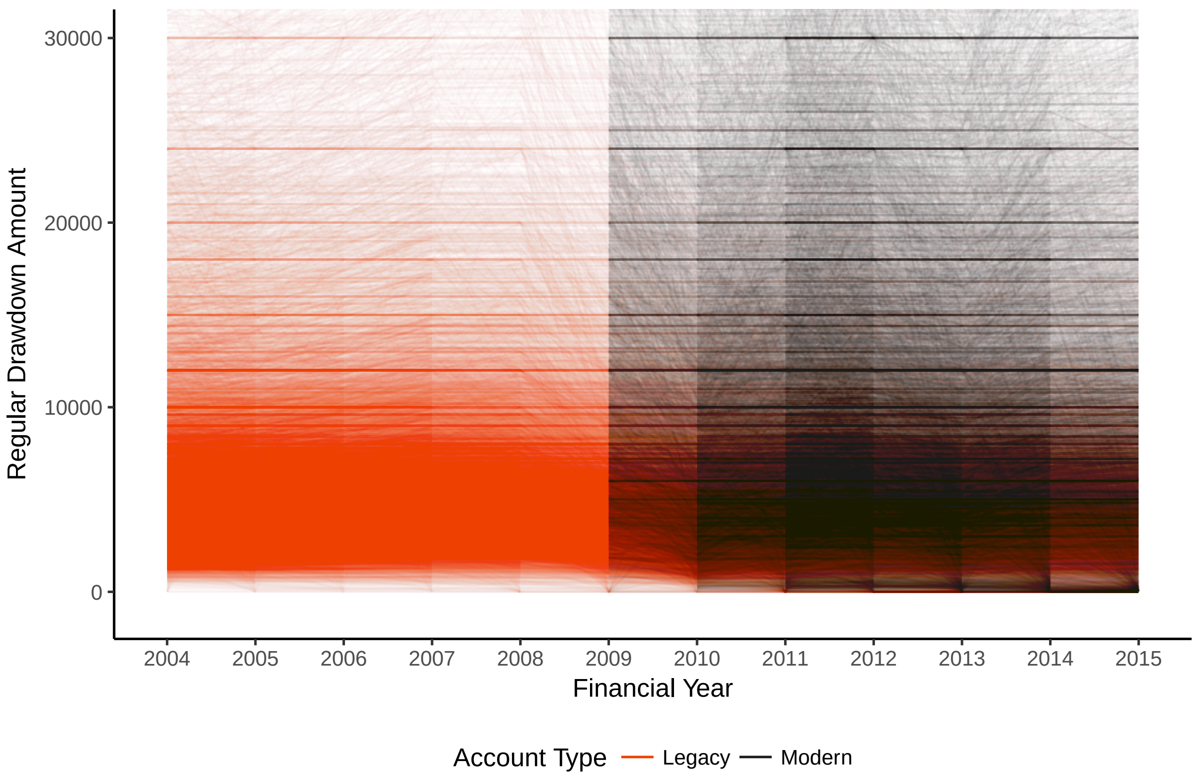

Figure 4.20: Account Balance at Financial Year Start Panel Visualisation

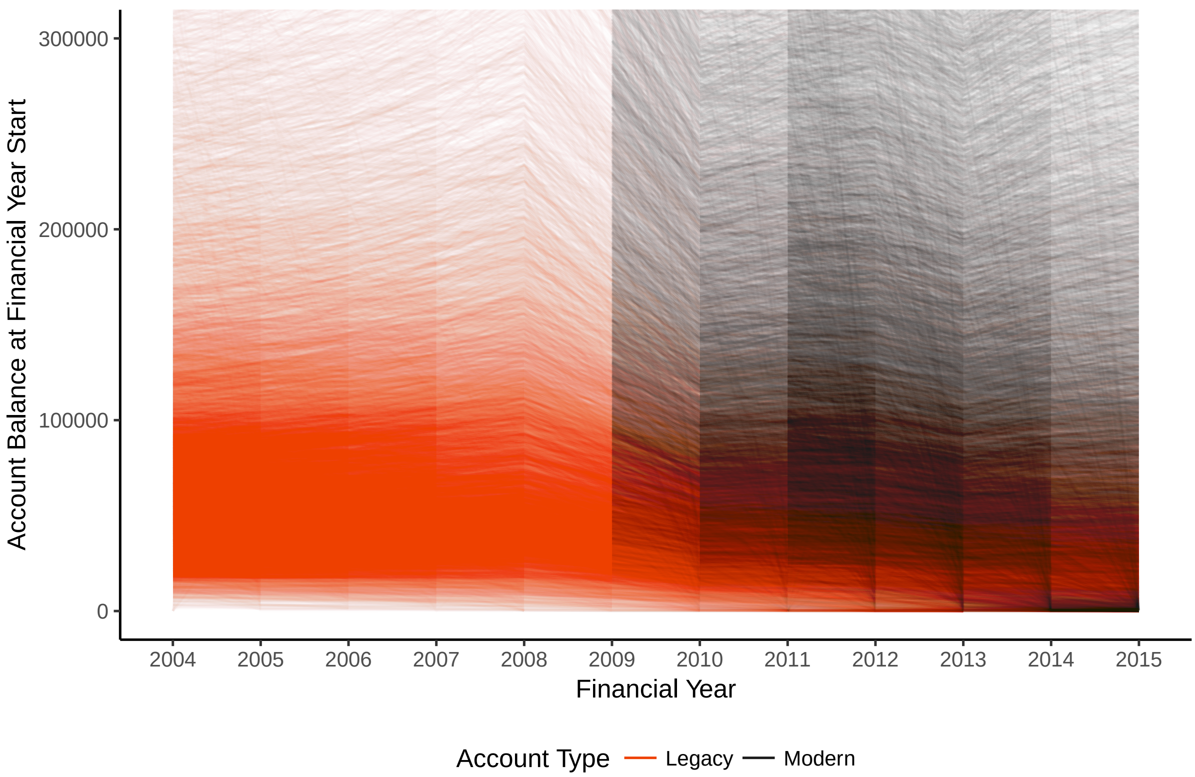

4.2.2 Manual Grouping

The panel visualisations suggest five groups that we can capture directly by applying filters to the underlying data. In section 4.2.4, we provide a sense of how large each of these groups are, both in terms of the number of retirees captured and in the proportion of the total sample this represents.

First are individuals who gravitated towards their respective minimum drawdown rates for all or most of the sampled time periods—shown in Figure 4.21. We allocated individuals to this group even if they lagged a year in adjusting to changes in the minimum drawdown rules or concessional rates, or if they strayed for one year briefly but otherwise faithfully followed this strategy.

Second, in Figure 4.22 we found a group of people that appeared to use the minimum drawdown rates as a strategy, but did not revise down their drawdown rates in years where concessional minima applied—financial years ended 2009–13. Again, some of these individuals lagged in adjusting to the new rules applying from 1 July 2007, and some drew at higher rates in some years but quickly returned to the non-concessional minima. Others seemed to be aware of the concessional minima, evidenced by their drawing below the non-concessional minima in some years—corresponding to rates below the zero line in this figure. These retirees seemed to prefer the non-concessional arrangements, however, and quickly returned to these rates.

Third, we find a group of retirees drawing regularly at a rate of 10% of their account balance annually, shown in Figure 4.23. 48% of this group was comprised of members with a TRIP, within which the maximum allowable drawdown rate is 10%.

A fourth group, visualised in Figure 4.24, are those who have a strong tendency to draw level income streams—except when they occasionally revise this level amount up or down.

Finally, after allocating the previous four groups, a very small number of the remaining retirees show a tendency to draw the same annual rate from their account for several successive years. Figure 4.25 shows this group.

Thus apart from the case of drawing the minimum drawdown rates, which can stay constant for several years in succession depending on the financial year and the age of the retiree, the tendency to draw at constant rates is exceedingly rare to observe in practice.

4.2.3 Machine-Assisted Grouping

After manually identifying the previous five groups, we apply a hierarchical clustering methodology to classify the remaining individuals. On experimenting with different distance metrics and linkage methods, the most successful combination proved to be the Euclidean distance combined with the `R` implementation of Ward's linkage method (for details, see Ward Jr, 1963; Murtagh and Legendre, 2014). We performed clustering on the observed values of several dependent variables, including:

Figure 4.21: Manual Grouping – Follow Minima

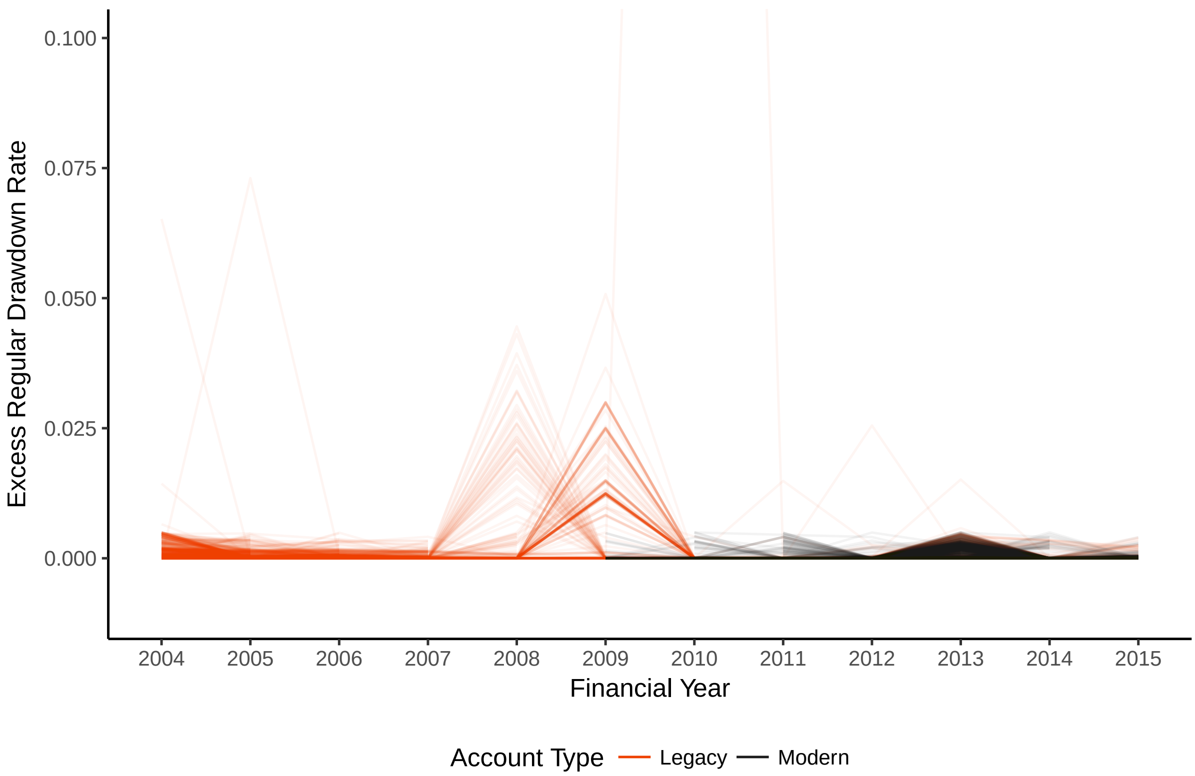

Figure 4.22: Manual Grouping – Follow Non-Concessional Minima

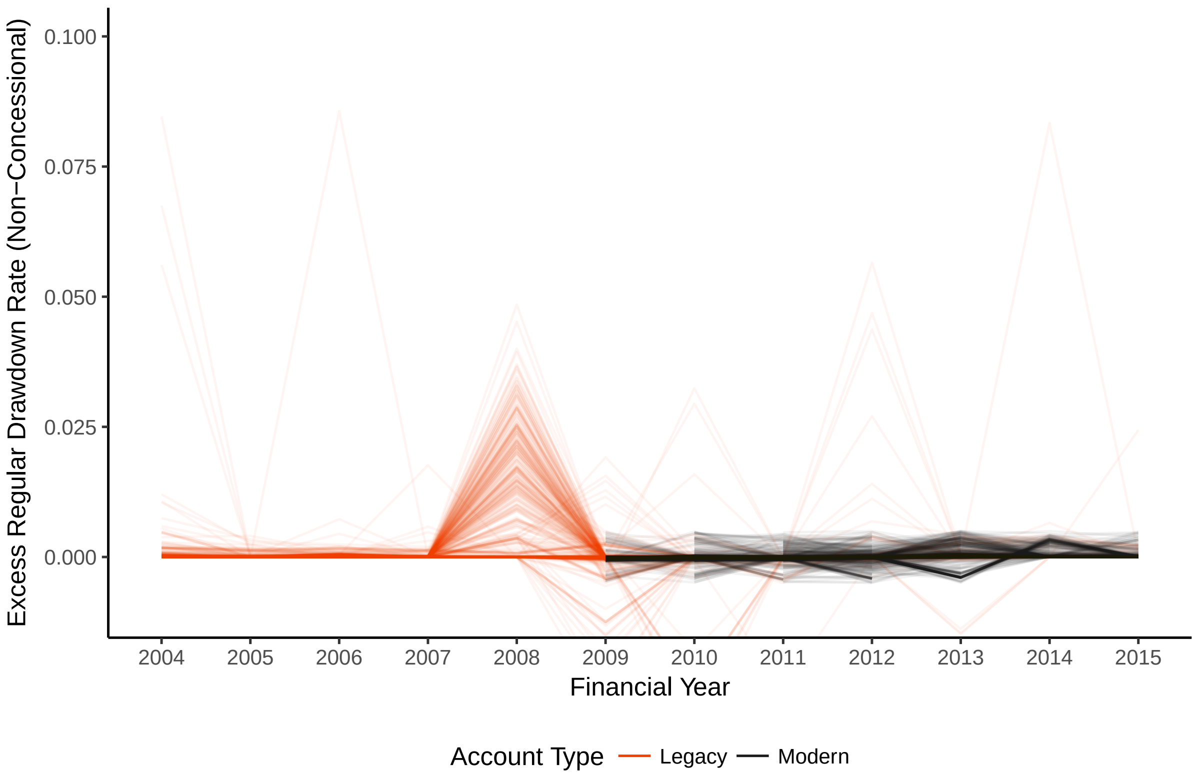

Figure 4.23: Manual Grouping – Draw 10%

Figure 4.24: Manual Grouping – Prefer Level Amount

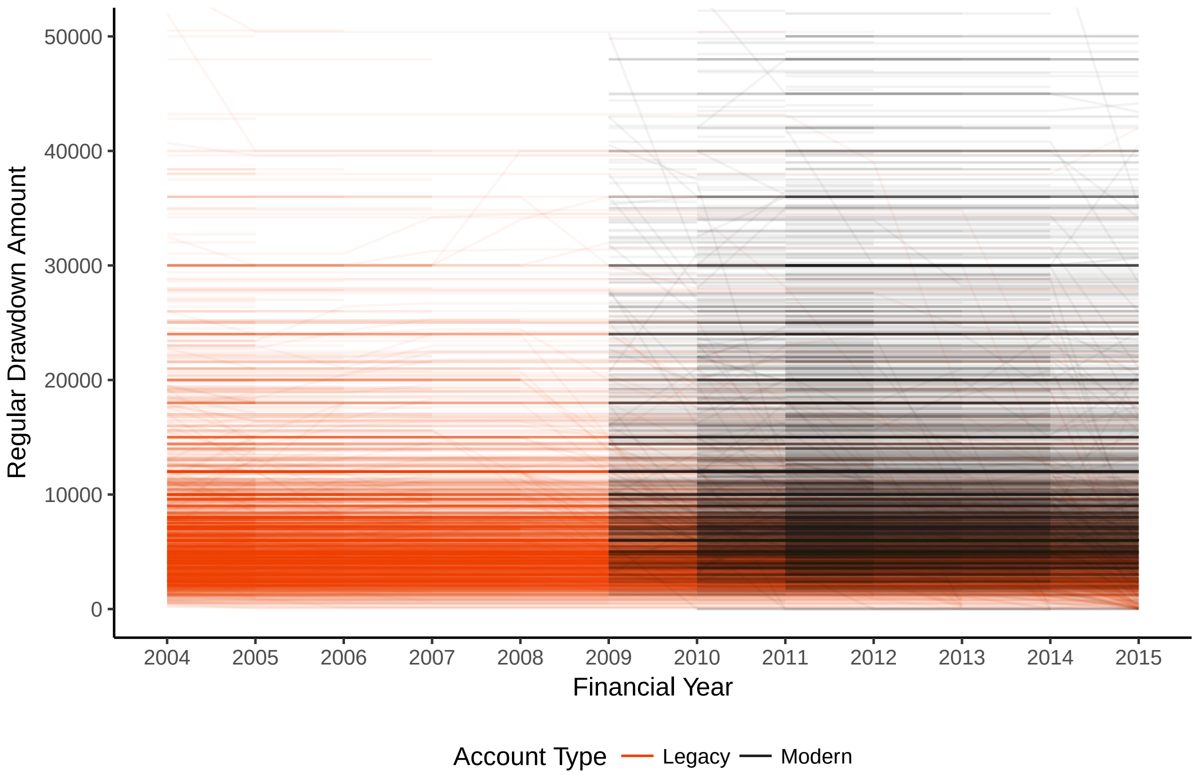

Figure 4.25: Manual Grouping – Prefer Level Rate

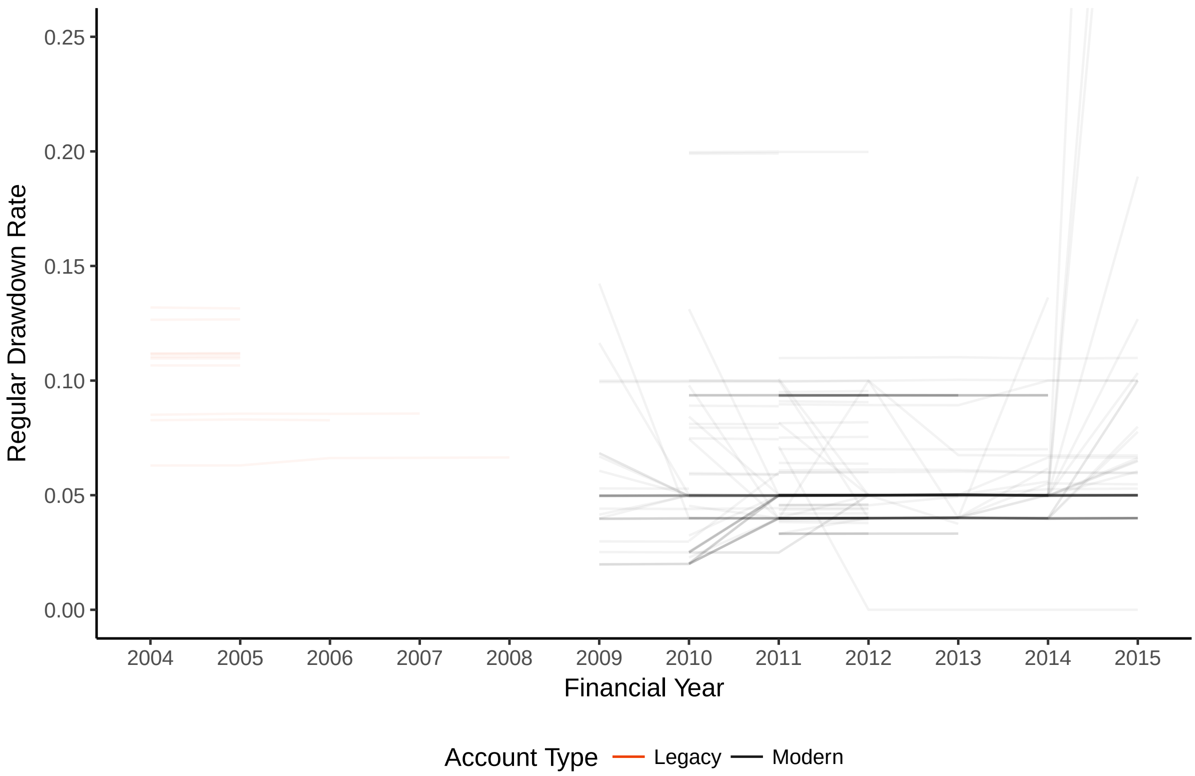

- The excess regular drawdown rate, ignoring the concessional rates applying in financial years ended 30 June 2009–13
- The excess regular drawdown rate, accounting for the concessional rates
- The first difference of the regular drawdown dollar amount
- The total drawdown rate, inclusive of regular and adhoc drawdowns

After the hierarchical clustering procedure grouped individuals who behaved similarly to each other into clusters, we inspected the results to determine against which we could attribute behavioural explanations. These individuals were allocated into the resulting 'clusters'.

Figure 4.26 shows multiple panels for individuals who used a combination of the concessional and non-concessional minimum drawdown rates as a guide. The vertical axis is the excess regular drawdown rate, ignoring concessional rates. As before, we allowed individuals a grace period around the time the rules changed, if they subsequently displayed a strong tendency to use the minima.

Figure 4.27 portrays a group primarily focused on the non-concessional minimum drawdown rates. Some of these individuals are not necessarily distinguishable from the previous cluster, however for the final allocation we aggregate these two clusters into the same behavioural group driven by one heuristic.

Figure 4.28 finds another group of retirees that try to draw at the minimum drawdown rates for most observed periods. These differ from those found using the manual rules in that they

Figure 4.26: Machine-Assisted Grouping – Use Concessional and Non-Concessional Minima

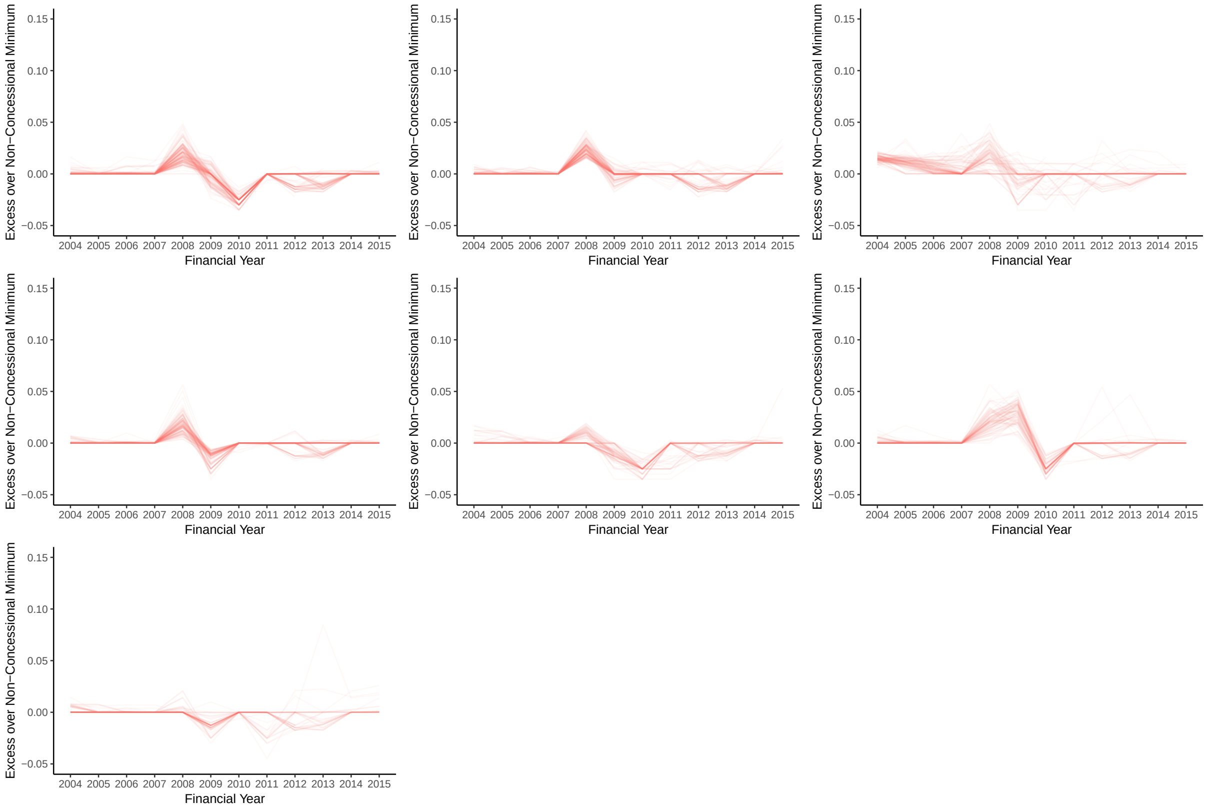

Figure 4.27: Machine-Assisted Grouping – Use Non-Concessional Minima

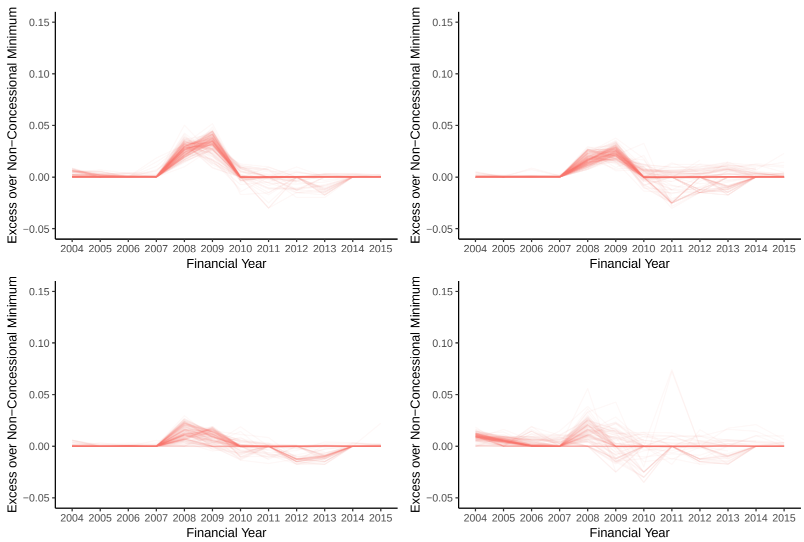

Figure 4.28: Machine-Assisted Grouping – Use Concessional Minima

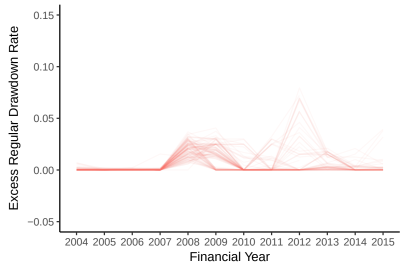

struggled longer and harder with the change in rules or the impact of the GFC. However they still exhibit the same tendency to follow the minimum rates, especially prior to the 2008 financial year.

Another group that seemed intent on drawing at or very close to the minimum rates prior to 2008 is shown in Figure 4.29. These retirees were unable to recover after the GFC as quickly as others who followed the minimum rules. In part, this may be due to a large reduction in account balances over financial years 2008 and 2009, leaving them with a much smaller denominator from which to form the drawdown rate.

Many retirees held a level drawdown amount over time. Figure 4.30 plots the first difference of the regular drawdown amount in dollar terms, such that the zero line represents a level income stream.

Another common, related behaviour, seen in Figure 4.31, was to draw the same dollar amount for most of the observed periods, but revising down the level amount at one stage. The dips correspond to the years in which retirees reduced the amount of their level income stream, and subsequently held the drawdowns level at this lower amount.

Figure 4.32 shows the less common, inverted behaviour: a level income stream with an upwards revision.

Similar to the group of individuals who were able to maintain the minimum drawdown rates until financial year 2008, we found a group which was able to draw a level amount until at least the 2007 financial year, but thereafter lost the ability or desire to hold a constant income stream. These retirees are visualised in Figure 4.33.

Finally, after allocating individuals into the above clusters, we identified a group of people who completely drew down their account balances while under observation. We see these retirees in Figure 4.34, where the total drawdown rate is on the vertical axis. Dropping down to a zero—or near-zero—drawdown rate after a near-complete liquidation of the account balance is possible since minimum drawdowns are only enforced if the dollar amount required exceeds \$10.

Figure 4.29: Machine-Assisted Grouping – Follow Minima 2004–7

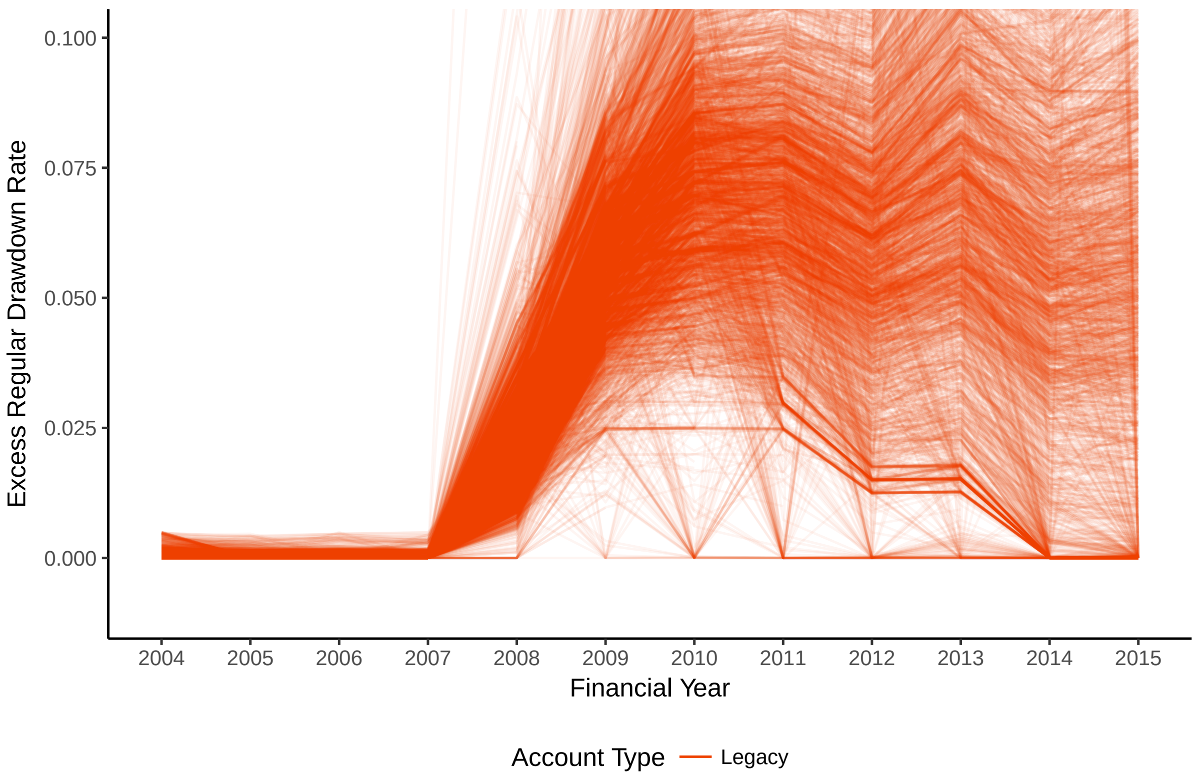

Figure 4.30: Machine-Assisted Grouping – Prefer Level Amount

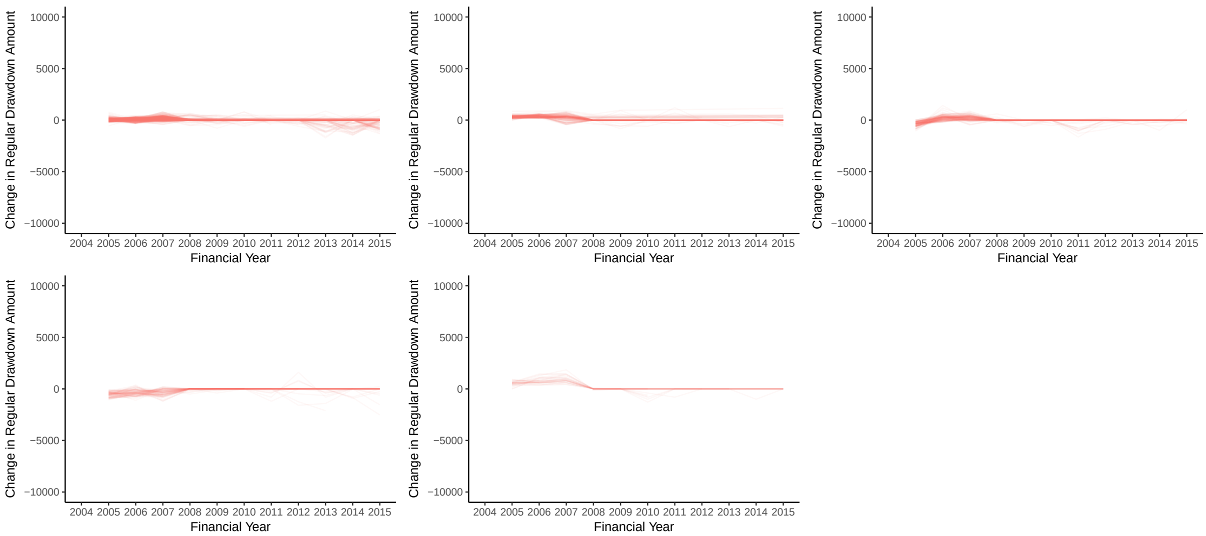

Figure 4.31: Machine-Assisted Grouping – Level Amount with Step Down

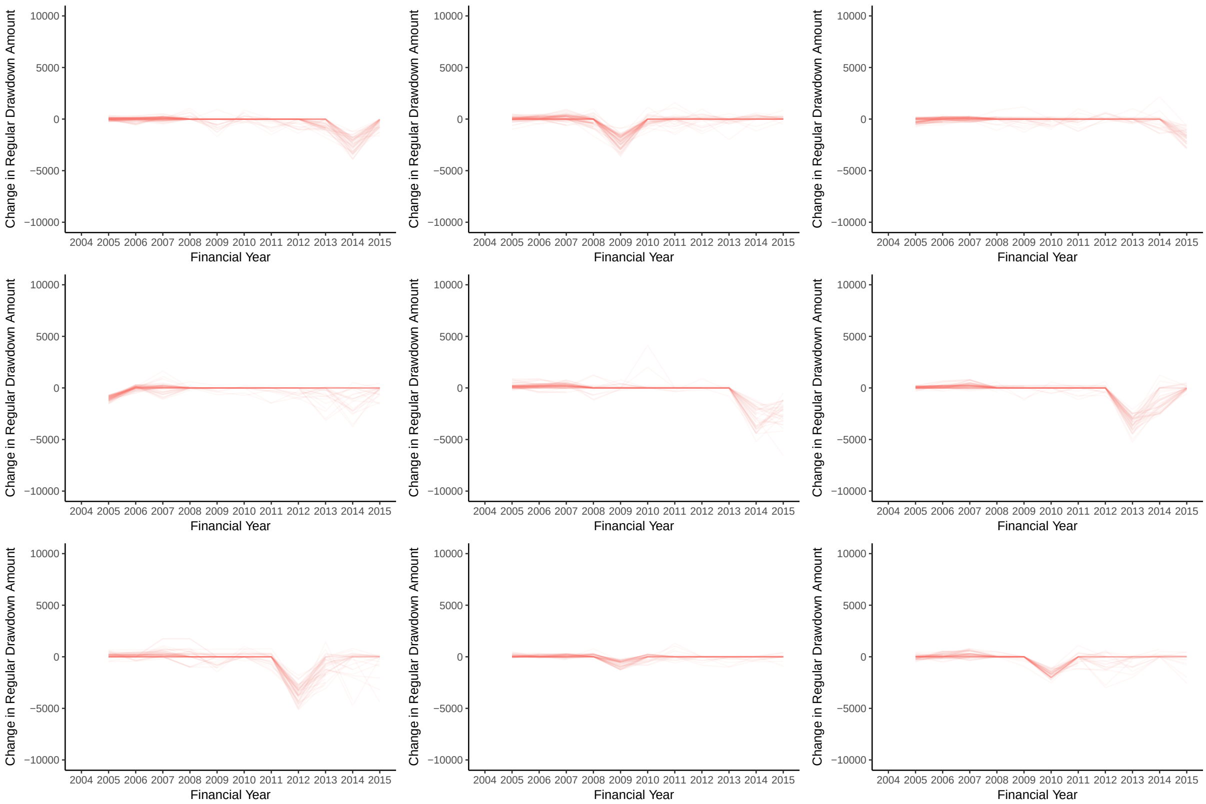

Figure 4.32: Machine-Assisted Grouping – Level Amount with Step Up

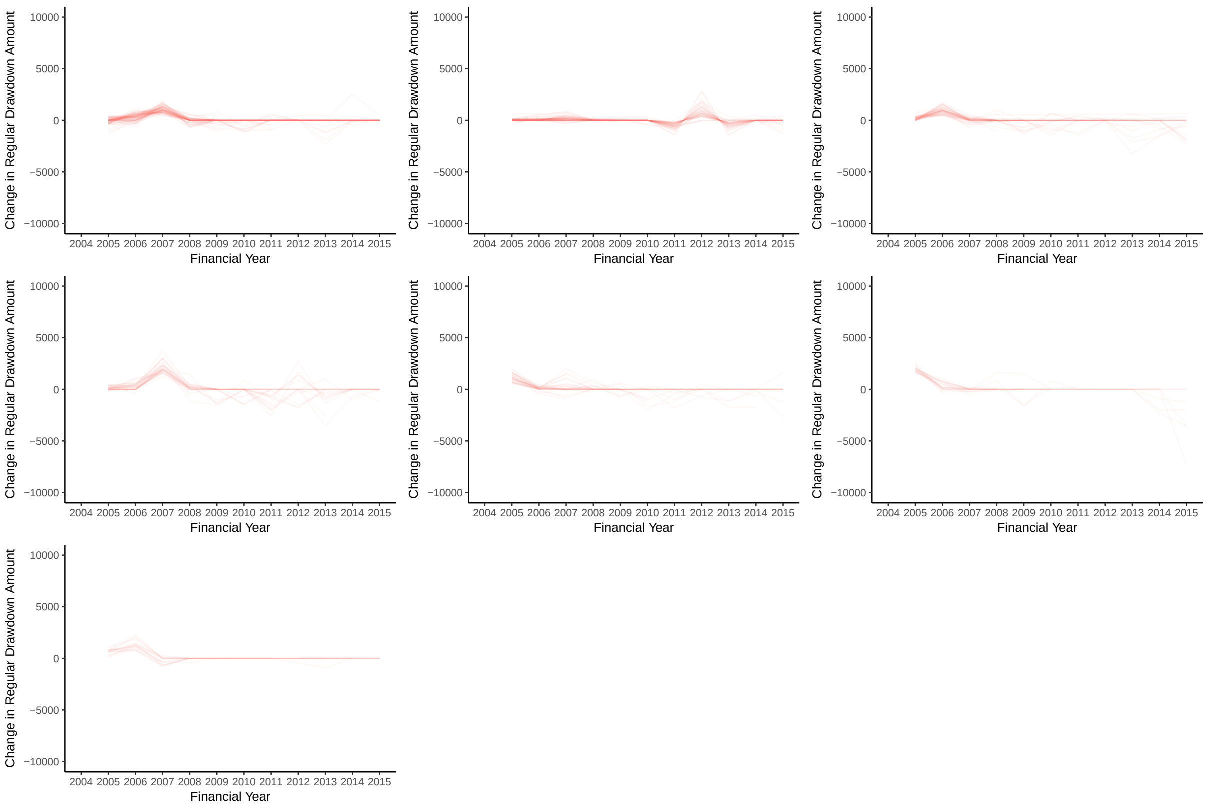

Figure 4.33: Machine-Assisted Grouping – Level Amount 2004–7

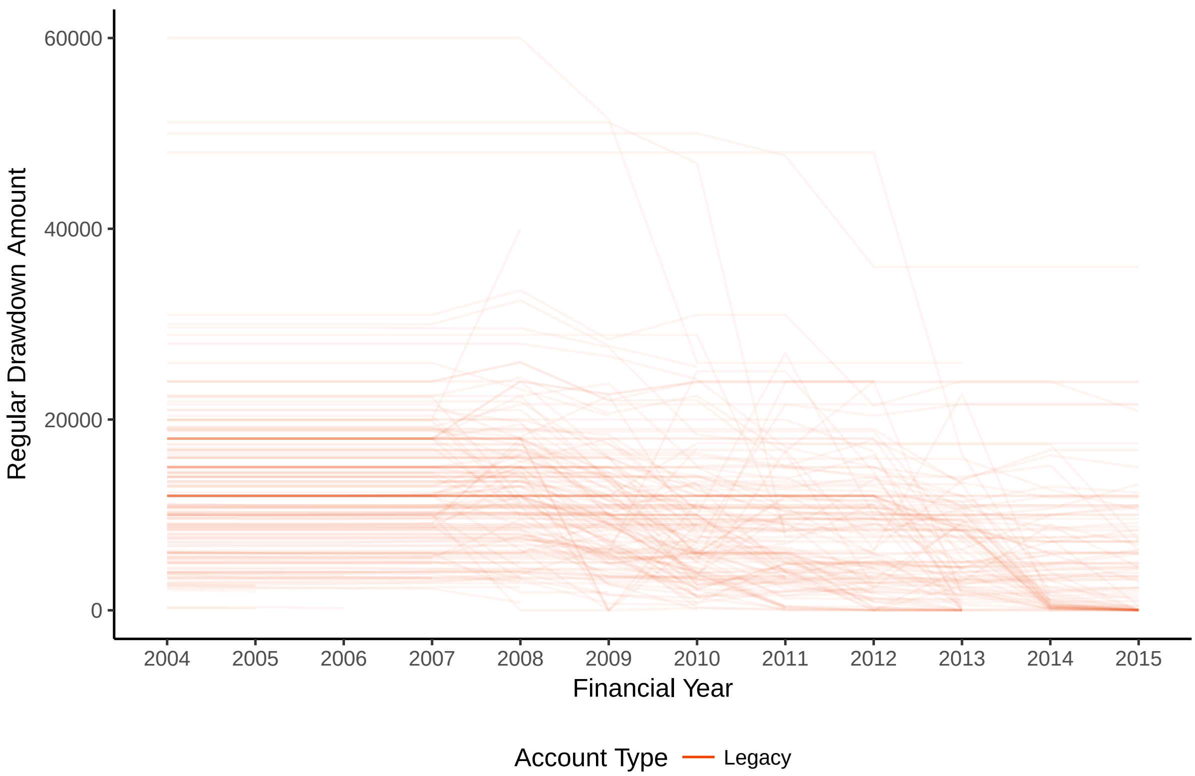

Figure 4.34: Manual Grouping – Complete Account Drawdown In-Sample

4.2.4 Final Cluster Allocation

After completing the manual and machine-assisted clustering procedure, we obtain the cluster allocation given in Table 4.21, where each cluster has its own unique economic interpretation. 17% of the sample remains unallocated to a discernible drawdown strategy.

However, we suggest that several of these clusters relate to identical heuristics, despite some variation in the execution. For example, for all clusters where the minimum drawdown rates—either concessional or non-concessional—are used as the chosen drawdown rates for multiple successive years, we attribute the same heuristic to describe the behaviour: using the minimum drawdown rates as a guide.

On aggregating these economically similar clusters into 'cluster groups', we arrive at the cluster group allocation shown in Table 4.22. Perhaps unsurprisingly, almost half of our (large) sample defaults to following the legislated minimum drawdown rates for a significant proportion of the observation period. Furthermore, more than a quarter of the sample prefers to draw the same amount in consecutive years—except for instances in which they choose to revise the level of their constant income stream. Such revisions are seldom in pursuit of a higher drawdown amount.

Importantly, none of these individuals appear to be protecting their regular income streams from inflation—in fact, using our methodology, we did not find evidence for inflation-adjusting behaviour at all. This is not necessarily indicative of diminishing purchasing power of retirees, as we are only able to observe the portion of their retirement income derived from an account-based pension. Other possible sources of income, such as their Age Pension entitlement, or investment income originating outside of their account-based pensions, may naturally grow at least as fast as inflation. Despite this, superannuation funds and financial advisors might be able to provide a better service to retirees by assisting them to draw inflation-adjusted income streams.

A small portion (4%) of our sample corresponds to younger retirees using transition to retirement accounts at their maximum allowable rate of drawdown: 10% per annum. Furthermore, after removing all other explicable behaviours, another 4% of the sample completely drew down their account balance while under observation.

4.3 Component 3: Categorical Regression Modelling

After constructing the cluster group allocation, the sample is split into four behavioural groups, as well as an unallocated group—representing 'noise'.

We report summary statistics for the time-invariant regressors in Table 4.23. To investigate how these groups differ statistically in the available time-invariant regressors, we perform a categorical regression using the multinomial logistic model. Table 4.24 summarises the regression output.

Table 4.21: Final Cluster Allocation Table

Cluster ID	Cluster Name	Cluster Size	Proportion of Sample
1	At Minima	7236	0.17
2	At Minima ('04–07)	4895	0.11
3	At Non-Concessional Minima	6891	0.16
4	10%	1811	0.04
5	Quickdraw	1549	0.04
6	Level Amount	6784	0.15
7	Level Amount ('04–07)	187	0.00
8	Level + Step Down	4331	0.10
9	Level + Step Up	715	0.02
10	Level Rate	172	0.00
11	Below Non-Concessional Minima	1786	0.04
12	Unallocated	7438	0.17

Table 4.22: Final Cluster Group Allocation Table

Cluster Group ID	Cluster Group	Cluster Size	Proportion of Sample
1	10%	1811	0.04
2	Quickdraw	1549	0.04
3	Follow Minima	20808	0.48
4	Level	12189	0.28
5	Unallocated	7438	0.17

Table 4.23: Summary Statistics for Candidate Regressors – Categorical Modelling

Variable	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Account Balance (First Year)	$0	$53,932	$95,541	$149,341	$181,114	$4,537,708
Risk Appetite	0.00	0.29	0.49	0.48	0.66	2.00
Age at Account Open	40.2	60.3	63.8	63.6	66.0	89.4
Age at 31 December 2015	57.8	70.9	75.8	75.6	80.5	103.8
Gender = Male	0	0	1	0.571	1	1
Legacy Account	0	1	1	0.534	1	1

Table 4.24: Cluster Group Allocation – Multinomial Logit Regression Model Output

	Multinomial Logit Model	
10%		
Gender = Male	0.719***	(0.0640)
Log Account Balance (First Year)	−0.637***	(0.0387)
Age at Account Open	−0.213***	(0.00747)
Risk Appetite	0.767***	(0.127)
Legacy Account	−2.521***	(0.0953)
Constant	18.25***	(0.720)
Quickdraw		
Gender = Male	0.297***	(0.0630)
Log Account Balance (First Year)	0.272***	(0.0406)
Age at Account Open	−0.0377***	(0.00633)
Risk Appetite	−0.414**	(0.147)
Legacy Account	−0.957***	(0.0844)
Constant	−2.889***	(0.679)
Follow Minima (base outcome)		
Gender = Male	0	(.)
Log Account Balance (First Year)	0	(.)
Age at Account Open	0	(.)
Risk Appetite	0	(.)
Legacy Account	0	(.)
Constant	0	(.)
Level		
Gender = Male	0.194***	(0.0270)
Log Account Balance (First Year)	−0.371***	(0.0175)
Age at Account Open	−0.0286***	(0.00282)
Risk Appetite	0.231***	(0.0591)
Legacy Account	−0.189***	(0.0331)
Constant	5.484***	(0.291)
Unallocated		
Gender = Male	0.304***	(0.0340)
Log Account Balance (First Year)	0.374***	(0.0222)
Age at Account Open	−0.0433***	(0.00357)
Risk Appetite	0.583***	(0.0753)
Legacy Account	−0.553***	(0.0445)
Constant	−2.902***	(0.371)
Pseudo-R^2	0.0581	
Observations	32280	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The base group is those retirees who closely follow the legislated minimum drawdown rates—which represents the default option in the fund. The estimated coefficients for the remaining four groups represent changes to the log odds ratio—relative changes to the odds ratio—of being in the respective group, relative to the base group, for unit increases in the available regressors.

As with the binary logistic regression models from section 4.1.2, the regression output does not translate directly into the change in probability of belonging to a selected group. Furthermore, in this nonlinear model, this change in probability relative to a unit change in a regressor also depends on the level of all the other regressors. Thus, as in the binary choice models, we will approximate the average marginal effects of each regressor on the total probability of belonging to each cluster.

The regression output table does, however, provide one especially interesting insight: how the two largest groups—those who follow the minima and those who maintain a level income stream—differ statistically in the available covariates. By observing the signs and statistical significance of the coefficients for the latter group, we can conclude that: males are more likely to be found in the group drawing a constant dollar amount; those with larger account balances are more likely to be in the group following the minima; delaying retirement increases the probability of following the minima; riskier investment allocations increase the probability of drawing a constant amount; and accounts opened before the current drawdown rules came into effect on 1 July 2007 were more likely to follow the minimum rates.

To examine the overall magnitude of the regressor effects, rather than the direction of the change as compared to a base case, Table 4.25 provides the average marginal effects. These are directly interpretable as the change in probability of belonging to a particular group relative to changes in the regressor values.

Most of these average marginal effects are quite modest in magnitude. For example, each year an individual delays retirement, the probability of following the minimum drawdown rates increases by about 1.1%, while a doubling of one's account balance in the first year of observation only increases the probability of consistently drawing at minima by approximately 3.0%.

Table 4.25: Cluster Group Allocation – Multinomial Logit Model Average Marginal Effects

	10%	Quickdraw	Follow Minima	Level	Unallocated
Gender = Male	0.0205***	0.00494*	−0.0644***	0.0156**	0.0234***
	(0.00220)	(0.00209)	(0.00554)	(0.00505)	(0.00409)
Log Account Balance (First Year)	−0.0218***	0.0111***	0.0297***	−0.0853***	0.0662***
	(0.00130)	(0.00132)	(0.00356)	(0.00311)	(0.00261)
Age at Account Open	−0.00683***	−0.000356	0.0109***	−0.00123*	−0.00246***
	(0.000269)	(0.000201)	(0.000570)	(0.000513)	(0.000415)
Risk Appetite	0.0208***	−0.0215***	−0.0790***	0.0176	0.0621***
	(0.00432)	(0.00492)	(0.0124)	(0.0110)	(0.00904)
Legacy Account	−0.0812***	−0.0226***	0.119***	0.0218***	−0.0370***
	(0.00344)	(0.00279)	(0.00677)	(0.00606)	(0.00531)
Observations			32280		

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Legacy accounts, however, were on average 12% more likely to follow the minimum drawdown rates for an extended period of time, after controlling for the other regressors.

Finally, we examine some model diagnostics for overall fit. One indicator is the Pseudo-R^2 of the model, which is low at 5.8%. To investigate the poor fit, in Table 4.26 we derive a multinomial extension of the classification table used to assess the overall explanatory power of binary choice models. In the multinomial case, the decision rule for predicting cluster allocation was to place the individual in the cluster which had the highest predicted probability of adherence, across the five possible outcomes.

This table, referred to as the confusion matrix, shows poor classification ability. Although the model correctly classifies approximately 93% (14092/15192) of the individuals in the group using the minimum drawdown rates into this group (sensitivity), only 48% (14092/29416) of the total predictions for an individual belonging to this group are accurate (positive predictive power—PPV). For the group which tended towards level income streams, the sensitivity is only 6.5% (617/9435) and the PPV is 43% (617/1438). The group who drew through their entire account balances while under observation, representing the 4% of the sample, had no individuals allocated to it by the model.

Consequently, we are confident that the available administrative data does not capture the majority of the variation in the observed cluster allocations. To further explore the reasons which may drive an individual to belong to a particular behavioural group, this area of the literature will either need to collect a richer set of demographic data on the individuals, or turn to studies which directly survey individuals to find reasons for their behavioural responses.

4.4 Other Results and Illustrations

In this section we investigate other data-driven insights into how our sampled retirees utilise their flexible account-based pensions. In particular, product designers should consider these important results when designing more appropriate income products tailored to this group of retirees.

Table 4.26: Cluster Group Allocation – Multinomial Logit Model Confusion Matrix

	Observed					
Predicted	10%	Quickdraw	Follow Minima	Level	Unallocated	Total
10%	145	12	150	77	69	453
Quickdraw	0	0	0	0	0	0
Follow Minima	985	1,046	14,092	8,606	4,687	29,416
Level	105	28	565	617	123	1,438
Unallocated	77	59	385	135	317	973
Total	1,312	1,145	15,192	9,435	5,196	32,280

4.4.1 Comparing Regular and Adhoc Drawdown Utilisation Rates

Broadly, 70% of all dollars drawn from account based pensions in our sample were derived from regular drawdowns, while the remaining 30% is attributed to the adhocs. Table 4.27 provides the breakdown.

Within the cluster group seeking to draw constant regular amounts through time, this proportion differed. As seen in Table 4.28, for this large group, covering 28% of all retirees observed, the respective allocation to regular and adhoc drawdowns is 82% and 18%. For all other retirees observed, the breakdown is seen in Table 4.29, where the allocation is approximately 66–34%.

We conclude that retirees who decide to use their accounts to provide a level income stream throughout retirement—with the level amount possibly revised after commencing—use less of their account balance on adhoc drawdowns, compared to the rest of the sample. These individuals—self-annuitisers—display a stronger desire for a relatively constant income stream, and rely less on their account balances for lump-sum withdrawals.

Modest Self-Annuitisation

We investigate further the level income streams generated by these self-annuitisers in Figure 4.35 and Table 4.30. Together, these illustrate that 50% of this group are generating level income streams of less than $5800. This is a surprising and profound discovery, as it voids one possible reason offered for why Australia observes very low levels of lifetime annuity sales: that there is no demand for modest income streams. As at 16 October 2017, Challenger offered 65-year-old females a guaranteed lifetime nominal income stream of roughly $7000 per year in exchange for a $100,000 up-front payment, and approximately $7400 for 65-year-old males (Challenger Limited, 2017b). Figure 4.36 and Table 4.31 show that for the 1171 retirees who are both in the level drawdown amount group and aged 65 in their first year of observation, approximately 50% have a balance between $54,000 and $154,000 in this year of age.

In fact, for this middle 50%, the equivalent guaranteed lifetime annuities Challenger could provide at a rate of 7% ($3780 and $10,780) correspond closely to the middle 50% of the level drawdown amount distribution in Table 4.30 ($3400 and $9500). Clearly, the level amounts that insurers can guarantee for life are of comparable magnitudes to those which retirees in account-based pensions already generate for themselves. Moreover, the up-front costs of these guaranteed annuities are of equally comparable magnitude to the account balances these individuals use to generate their own income streams.

Consequently, this empirical data analysis does not support the argument that retirees avoid

Table 4.27: Aggregate Regular and Adhoc Drawdown Breakdown – Entire Sample

Aggregate Regular Drawdowns	Aggregate Adhoc Drawdowns	Proportion Regular	Proportion Adhoc
$2,280,307,703	$970,617,135	0.70	0.30

Table 4.28: Aggregate Regular and Adhoc Drawdown Breakdown – 'Level' Cluster Group Only

Aggregate Regular Drawdowns	Aggregate Adhoc Drawdowns	Proportion Regular	Proportion Adhoc
$688,147,775	$155,750,943	0.82	0.18

Table 4.29: Aggregate Regular and Adhoc Drawdown Breakdown – Excluding 'Level' Cluster Group

Aggregate Regular Drawdowns	Aggregate Adhoc Drawdowns	Proportion Regular	Proportion Adhoc
$1,592,159,928	$814,866,192	0.66	0.34

Figure 4.35: Histogram of Regular Drawdown Amount – 'Level' Cluster Group Only

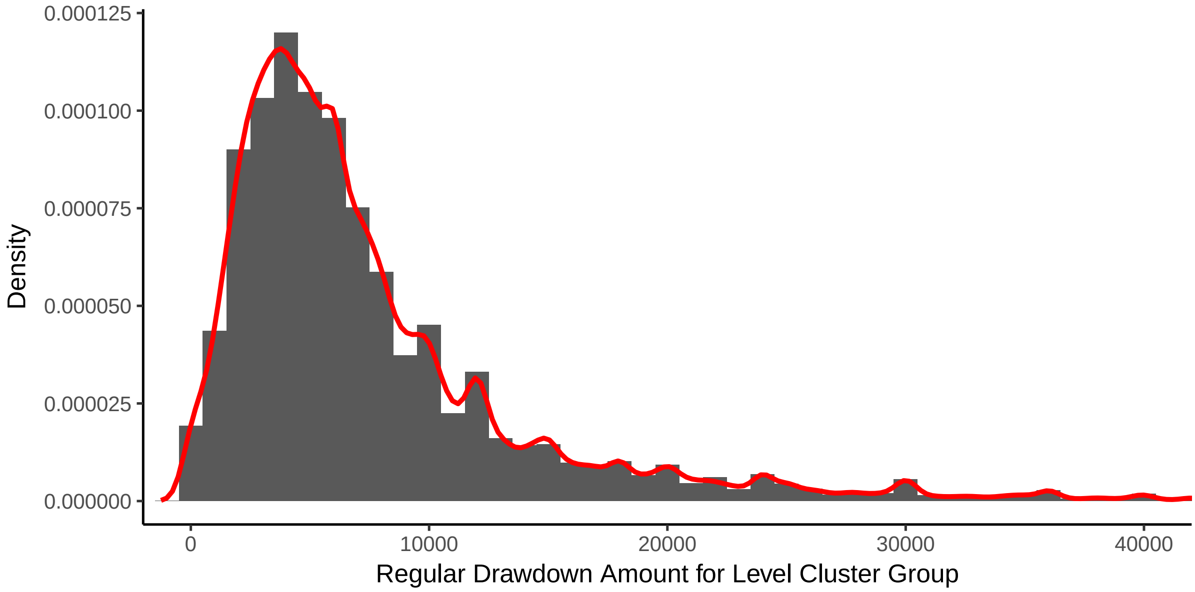

Table 4.30: Summary Statistics for Regular Drawdown Amount – 'Level' Cluster Group Only

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Regular Drawdown Amount	0	$3444	$5760	$7823	$9492	$120,000

Table 4.31: Summary Statistics for Account Balance in First Observed Year – 'Level' Cluster Group and Over Age 65 Only

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Account Balance	0	$53,980	$90,097	$123,771	$154,000	$1,344,842

Figure 4.36: Histogram of Account Balance in First Observed Year – 'Level' Cluster Group and Over Age 65 Only

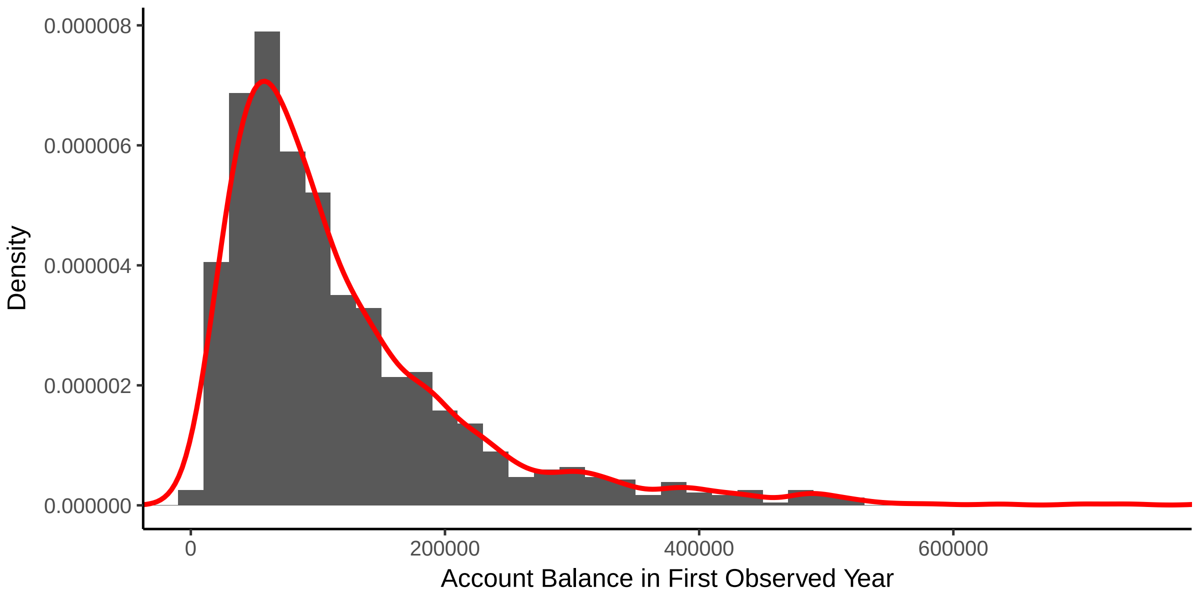

lifetime annuities due to a mismatch between their available superannuation assets and their desired income stream amounts—or that the rate of return on investment for lifetime annuities is too low. Instead, it is more likely that these individuals have used their superannuation assets to open an account-based pension due to other factors—such as the flexibility of future income amounts, and the ability to make adhoc withdrawals.

Adhoc Utilisation Rates

In our sample, 35% of individuals make at least one adhoc drawdown while under observation. By summing each individual's regular and adhoc drawdowns made while observed, we can calculate the proportion of their overall drawdown amounts attributed to adhocs—and call this the adhoc utilisation rate for each account. The distribution of this variable is summarised in Figure 4.37 and Table 4.32.

For those who make adhoc drawdowns, the distribution of the utilisation rate appears broadly uniform for most of the interval $(0, 1)$. Reaching a utilisation rate approaching 100% is rarer, and represents individuals who draw through their accounts quickly—relative to their nominated regular drawdowns—using adhocs.

Table 4.32: Summary Statistics for Adhoc Utilisation Rate

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Adhoc Utilisation Rate	0.00	0.20	0.45	0.46	0.73	1.00

Figure 4.37: Histogram of Adhoc Utilisation Rate

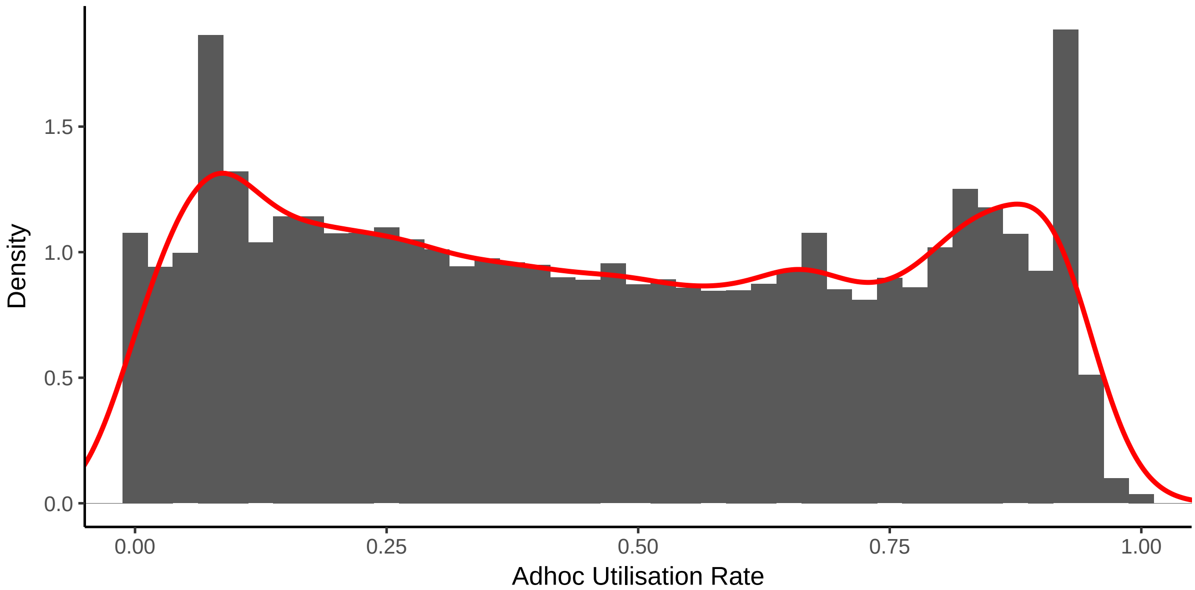

Adhoc Drawdown Strain

Finally, we look at the distribution of the adhoc drawdown amounts made during the observation period in Figure 4.38 and Table 4.33.

There is a visible tendency for retirees to draw adhoc amounts in multiples of $5000. As these nominal amounts do not provide a view into what strain these adhocs place on the remaining account balance, we manipulate the data to obtain this perspective. Only for individuals observed for at least 7 years and making at least one adhoc drawdown in this period, we average the proportion of account balance drawn as adhocs over this period to create a time-averaged adhoc drawdown rate—interpreted as the average annual strain placed on account balances to fund adhoc drawdowns. This adhoc drawdown strain variable has distribution as per Figure 4.39 and Table 4.34. For those who make adhoc drawdowns, the median level of adhoc drawdown strain contributed to a reduction in account balances over time at the rate of 4% per year on average.

Table 4.33: Summary Statistics for Adhoc Drawdown Amount

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Adhoc Drawdown Amount	$1	$3039	$9000	$29,583	$20,620	$2,730,756

Figure 4.38: Histogram of Adhoc Drawdown Amount

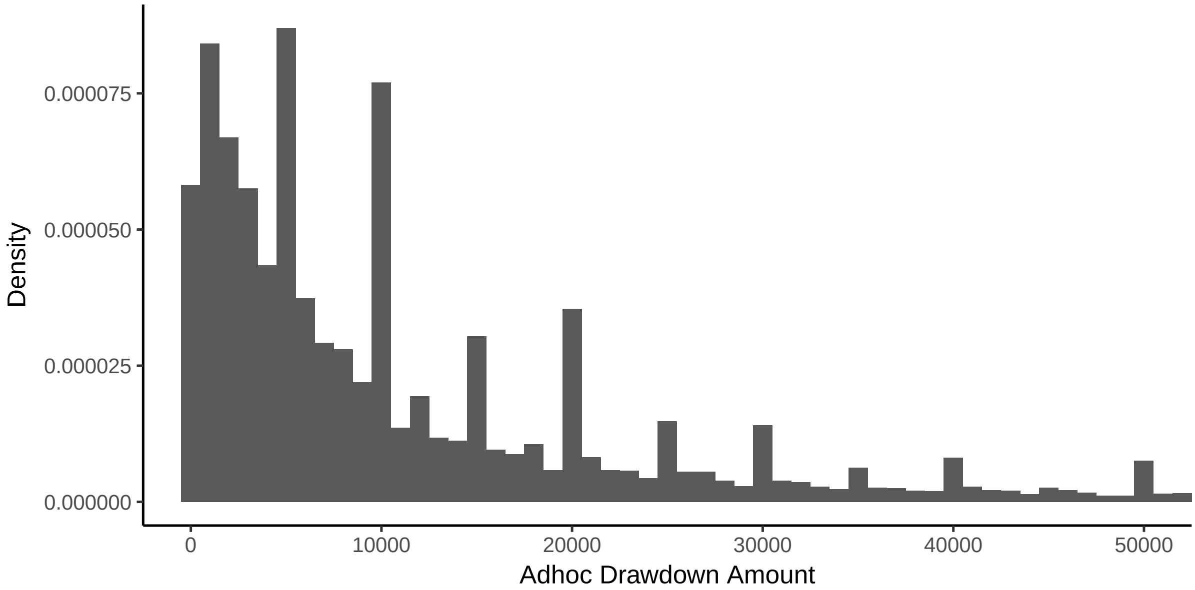

Table 4.34: Summary Statistics for Adhoc Drawdown Strain

	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
Adhoc Drawdown Strain	0.00	0.01	0.04	0.07	0.10	0.58

Figure 4.39: Histogram of Adhoc Drawdown Strain

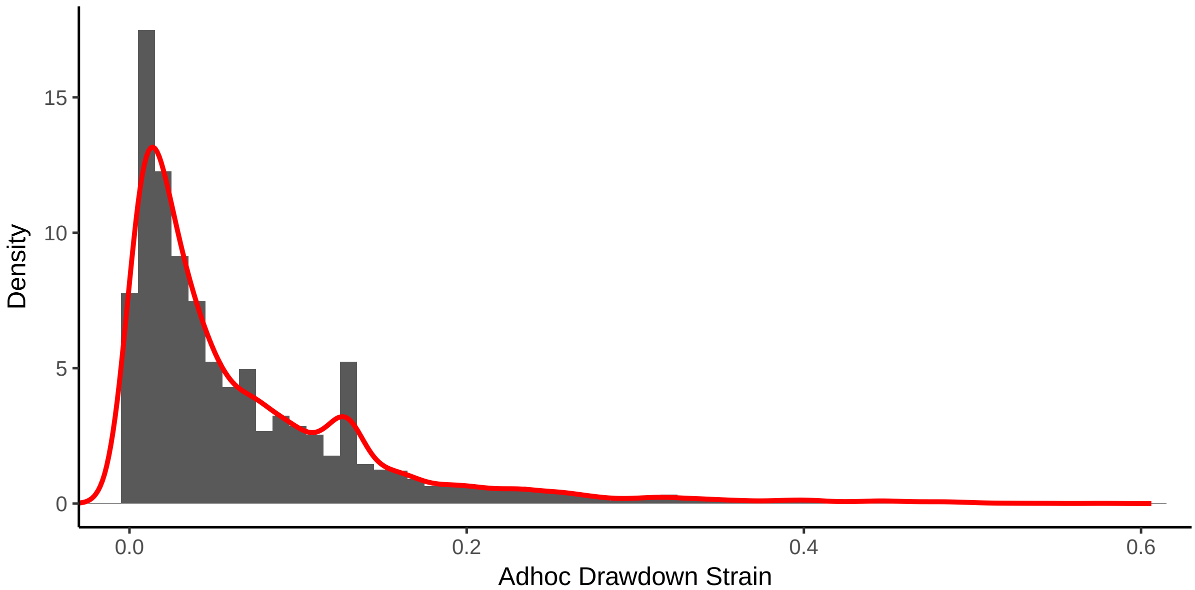

CHAPTER 5

DISCUSSION

The previous chapter presented and explored the results derived from the three components of the methodology. A series of panel regression models estimated the statistical relationship between the available characteristics and several dependent variables of interest. After fitting these models, limited explanatory power motivated a deeper analysis.

To introduce a behavioural dimension to the panel data, manual grouping and machine-assisted cluster analysis on the observed drawdowns through time found a small number of relatively large groups which appeared to follow very simple drawdown patterns.

With distinct behavioural patterns identified and qualitatively interpreted, a categorical regression model highlighted statistically significant differences in the distribution of characteristics displayed by members of these groups. In addition to the statistical components of the methodology, descriptive analysis provided other insights that policymakers, financial advisors and retirement income product designers can leverage in their work.

This chapter discusses the results and their implications within both academic and industry contexts.

5.1 Limitations

Unobserved Characteristics

Foremost, our dataset lacked several key variables that are not only interesting to study, but also prevalent in the retirement savings and decumulation literature. These include: marital—or spousal—status; health indicators; wealth not held in the retiree's phased withdrawal account; and income derived from other sources, such as the means-tested Age Pension and other investment income.

While panel modelling techniques can control for the unobserved, time-invariant characteristics—α_i—that influence drawdowns, and although we relied on asymptotic results in our large sample, at least two issues remain when missing these other key variables. First, no matter how robust our estimation techniques are, we cannot perform inference on the unobserved factors—and the effects of health and other wealth are of particular interest to the academic literature and the retirement incomes industry.

Second, the effects of the observed regressors are entangled with the effects of the unobserved, time-varying regressors—such as health and other wealth—that share correlation with both the dependent variables of interest and the observed regressors. For example, health may deteriorate with rising age, and therefore have a negative correlation with our observed age variable. Thus, when an age variable is statistically significant in determining drawdown rates, the partial effect of rising age is confounded with the partial effect of deteriorating health status. Although this limits prediction at the level of the individual, it does not limit the accuracy of the estimated effect of ageing as it applies to groups of retirees more broadly.

In the categorical regression model, a similar issue prevailed. Although able to find statistically significant regressors from the set of available characteristics, many unobserved factors might provide further insights into why retirees follow the observed behaviours.

Although not analysed in this project, the available dataset does contain information that might comment on the influence of deteriorating health on drawdown behaviours. A separate analysis on the individuals who died while under observation would uncover whether a 'proximity to death' variable significantly changes observed behaviours. Potentially, this could rival other imperfect health indicators, such as subjective health reported or healthcare expenditure. This remains for future work to determine.

One Asset Class

We reiterate that the current dataset on account-based pensions exists in isolation from information on members' other wealth and income. These factors likely also play a significant role in determining drawdown behaviours and rates. For example, when retirees can derive income from assets held outside of the superannuation system, they may draw on these with preference. In general, investment income earned on assets within the superannuation system is concessionally taxed. Moreover, during the observation period, for retirees aged 60 and over, investment earnings on superannuation assets incurred no tax. Consequently, drawing at the minimum retained more wealth within the low-tax superannuation environment.

This does imply that drawdowns at the minimum are not indicative of overall consumption patterns during retirement. However, drawdowns above the minimum are likely to represent consumption needs, assuming retirees only withdraw in excess of the minimum rates when they require the excess to fund consumption habits.

For the purposes of this project, which focused on the second pillar of Australia's retirement system, having observed only assets held in account-based pensions does not limit the relevance of the findings. The behaviours within these accounts were still identified and differentiated based on the observed characteristics, and subsequent sections will discuss the academic and social contributions of these findings.

Modelling Rates

The dependent variables of interest in this research were the rates at which retirees are drawing down from their phased withdrawal accounts. These have a useful interpretation in terms of the relative speeds of decumulation, something which modelling dollar amounts alone could not capture. However, a drawback to modelling rates—which are naturally constrained between 0 and 1—is that they are less likely to satisfy the assumptions of standard linear models.

As we demonstrated in Chapter 4, simple transformations allowed the rate variables to adhere more closely to the assumptions underlying these linear models. In other cases, where the dependent variable was discrete or contained significant probability masses, we employed nonlinear models, but at a cost—we lost the ability to directly interpret the coefficients on time-invariant characteristics, such as gender or our derived risk appetite metric.

Grouping Methodology

Broadly, the second component of our methodology—grouping individuals by observed behaviours over time—relied on classification using observed drawdowns, rather than directly surveying individuals. The behaviours identified by the manual grouping procedure are convincing, both visually and because they are inspired by economic reasoning and the theoretical literature.

In contrast, the machine-assisted cluster analysis results are less convincing. Hierarchical clustering is guaranteed to find as many groups in the dataset as desired, including patterns that may not necessarily represent members following a specific rule. Correspondingly, we only incorporated the results of the cluster analysis into the behavioural grouping if we could determine a clear rule underlying the trajectory of the drawdown rates displayed by each cluster. Consequently, it is possible that within the 'unallocated' cluster, there remain behavioural groups that evaded both manual and machine-assisted attempts at identification and classification.

5.2 Academic Contributions

5.2.1 Panel Modelling Contributions

As explored in Chapter 2, the empirical literature on behaviours within phased withdrawal accounts is underdeveloped—despite a theoretical body of literature exploring the optimal drawdown behaviours in these accounts. Until now, a lack of appropriate data was a limiting factor, but this paper shows that a panel dataset on mostly administrative variables can provide insights into the statistically and economically significant effects of characteristics such as age, gender and account balance on determining the rate of drawdown. Here, we summarise key findings from Chapter 4.

Members more likely to draw at the minimum rates:

- Are older
- Face higher minimum drawdown rates

Those more likely to make adhoc drawdowns:

- Are older
- Face lower minimum drawdown rates

Retirees drawing at higher rates tend to:

- Be male
- Be younger
- Have smaller account balances
- Have higher risk appetites
- Have retired older
- Be facing higher minimum drawdown rates

Members who put more strain on their account balances through adhoc drawdowns:

- Are older
- Have larger balances
- Face lower minimum drawdown rates

5.2.2 Behavioural Contributions

As well as better explaining drawdown rates over individual financial years, a second key advantage to having panel data is the ability to track individuals over time—identifying and distinguishing between observed behaviours.

A powerful finding, especially given our large sample size, is that two very simple rules explained the drawdown behaviours of more than three quarters of our sample. Almost one half (48%) of the observed retirees used the minimum drawdown rates as an anchor, while more than one quarter (28%) tended towards drawing level dollar amounts. Within this second group of retirees, at least 35% revised down the level of their income stream while observed—and others may still behave similarly in later, unobserved years.

Those following the minimum drawdown rates were more likely to:

- Be female
- Have larger account balances
- Have lower risk appetites

- Have retired later

By contrast, those drawing level amounts were more likely to:

- Be male
- Have smaller account balances
- Have higher risk appetites
- Have retired younger

These differences, while statistically significant, were relatively small in magnitude. However, members with accounts opened before the current minimum drawdown rates came into effect were, on average, 12% more likely to draw at the minimum, and 2% more likely to draw constant amounts—compared to their counterparts with newer accounts.

Two smaller groups collectively accounted for 8% of the sample. One of these—4% of the sample—was comprised of retirees drawing at or near 10% for all or most of their observed years. 88% of these accounts were TRIPs, which are subject to a maximum drawdown rate of 10% and can only be opened by retirees younger than 65. For the remaining 12% of this small behavioural group, the 10% rule may simply have been an attractive heuristic.

The second of the smaller groups corresponded to individuals drawing through their entire account balance while under observation. These individuals seemed uninterested in using their phased withdrawals to generate income for all or most of their retirement.

17% of our sample remained unallocated into any discernible behavioural group. While exploring this group in further detail is of economic interest, it remains for future work to investigate.

We can relate these findings directly to previous theoretical studies on optimal drawdowns, particularly the work of Bateman and Thorp (2008) reviewed in Chapter 2. One of the findings from their paper is that the legislated minimum drawdown rates—which came into effect on 1 July 2007—are a good guide to the simulated optimal drawdown pattern through retirement for a variety of assumed parameters in their calibrated utility functions. However, Bateman and Thorp showed that for some parameter values, a fixed drawdown rate heuristic provided higher income in earlier years of retirement, and increased utility relative to following the minima.

Observing that almost half of our sample used the minimum drawdown rates as a guide is consistent with these findings. Whether due to soft compulsion (default options), anchoring effects (fixating on numerical figures), sound financial advice, retiree introspection, or any other postulated reason, something is successfully driving individuals to follow the minima—which this literature finds is not far from the optimal behaviour.

By contrast, we did not find a prominent group of individuals attempting to draw through their account balance at a constant rate. After accounting for minimum drawdowns, the only evidence for constant drawdown rates was at the 10% level—and most of this was due to younger retirees maximising the value of a TRIP. However, the large group of individuals drawing constant *amounts*—at rates higher than the minima—suggest that many retirees behave consistently with results from the optimality literature—in that drawdown above the minimum is favourable during earlier years of retirement. Observing the common downwards revision of the level income streams generated by these individuals suggests that higher income is preferable in early retirement. This idea is consistent with research by Aegon (2016), finding that consumption levels are higher in the early years of retirement, but decrease by age 75.

Alternatively, Asher et al. (2017) and Hulley et al. (2013) find that retirees within the taper region of the Age Pension means test are more likely to decumulate their assets faster. It is generally preferable—from a tax-minimisation perspective—to retain as much money as possible in the tax-favourable superannuation system, decumulating other assets first. However, some individuals subject to the means test taper may have the majority of their wealth held in superannuation assets. Consequently, these members may be decumulating at higher rates earlier in retirement due to incentives introduced by the Age Pension means test.

As for not observing many retirees using a constant drawdown rate heuristic, our analysis of the data inspires the following suggested explanation: due to the volatility experienced by account balances over time, especially during economic downturns, drawing from these accounts at a constant proportion of the account balance may introduce undesirable volatility into a retiree's income streams. Research on consumption preferences suggests that realistic utility functions penalise volatility in incomes over time. For a discussion on the relationship between risk aversion and consumption smoothing over time, see Garcia et al. (2006). Our observations are consistent with the utility literature in that most retirees who deviate from the default minimum drawdown rates elect to draw constant dollar amounts.

Finally, it is worth noting that while drawing level dollar amounts over an extended period of time was a popular strategy, we did not observe a group of individuals protecting themselves from inflation by steadily increasing their drawdown amounts over time. That is, retirees are not protecting themselves from the erosion of spending power by inflation—at least, not through their account-based pensions. Over the period 2004 to 2015, the approximate cumulative effect of inflation was a 33% rise in the cost of living (derived from Australian Government Australian Bureau of Statistics, 2017). In this light, even retirees drawing level amounts for the entire observation period are, in real terms, decreasing their drawdown amount—although perhaps not intentionally. However, as Age Pension payment amounts are indexed to inflation, and we do not observe any income retirees derive from their other assets, we cannot comment on the overall erosion of their consumption power over time.

5.3 Social Implications

5.3.1 Policy

Minimum Drawdown Rates

Whether intended as soft-compulsion or simply as a conservative lower bound, the default option of drawing at the legislated minimum drawdown rates proved very popular amongst around half of the members in the analysed fund. This is consistent with findings from the behavioural economics literature, suggesting that default options strongly influence financial decisions relating to retirement incomes. As discussed in Chapter 2, Bateman et al. (2017) present a recent example.

As a result, we think it becomes clear that government decisions to change—or not change—the minimum drawdown rates do impact a large number of individuals throughout the duration of their retirement.

5.3.2 Retirement Income Product Design

On 1 July 2017, the Australian government relaxed the restrictive regulations that determine which retirement income products can retain favourable taxation treatment within the superannuation system—previously only afforded to traditional guaranteed lifetime and term annuities, and account-based pensions. Our findings suggest a suitable income product that insurers can now design for the Australian market, explored below.

Stepped Annuities

A nontrivial portion (10%) of our sample drew an income stream that resembled a 'stepped annuity'—an otherwise level income stream subject to a downwards revision in the level.

One interpretation for this behaviour is that members may desire a higher annual income earlier in retirement. Alternatively, members may revise down not because they want to spend less, but because they would like their income stream to last longer into the future.

In either case, the benefit of purchasing annuity products from a life insurer, as opposed to self-annuitising, is the longevity insurance an insurer can provide through guaranteed lifetime income. Beyond this, retirees may wish for annuity income to be higher initially, at the cost of reduced income later in retirement. The observed drawdown behaviours within account-based pensions suggest there may be demand for this. As a result, insurers should develop stepped annuities to be offered in the Australian market.

CIPR Options

An important finding from the analysis of the proportion of drawdown amounts attributed to adhocs is discovering the considerable heterogeneity in adhoc drawdown utilisation. At the aggregate level, adhocs account for 30% of the dollar amounts drawn down from account-based pensions. However, this single figure masks two key properties of the underlying distribution.

Foremost, we only observe 35% of our sample making at least one adhoc drawdown during the observation period—although this rate might become higher across a retirement time horizon of 20–30 (or more) years. Furthermore, within this group of members who make adhoc drawdowns, the adhoc-to-regular drawdown ratio is, roughly, uniformly distributed between 0 and 100%.

These results have implications for the development of CIPRs. As suggested by Treasury, funds could design CIPRs to provide both a longevity-protected income stream and an allowance for adhoc withdrawals throughout retirement (Australian Government The Treasury, 2016a). Prior to this research, a reasonable suggestion for presenting this option to members may have been to specify a default split of superannuation assets—a percentage which purchases a lifetime income stream, and the remainder placed in a liquid account. However, our analysis finds that the majority of retirees make no adhoc drawdowns over at least a 12-year period, while the remainder utilise this facility to highly variable extents. Consequently, no default option for this split would suit any more than, by chance, a handful of retirees.

Since we know defaults are powerful behavioural anchors, it may be more prudent to not specify a default split. Instead, funds should require members to make the decision with regard to their own expected needs in retirement, and provide financial advice to support them in this decision. CIPRs themselves could become default options for accessing accumulated wealth in DC accounts after retirement in Australia. In this eventuality, discussions with members as to their needs would be critical to ensure appropriate allocation of assets.

Finally, a barrier to the widespread appeal of these more advanced decumulation arrangements may be the inherent unpredictability members feel when considering a retirement time horizon of 30 years. For this reason, account-based pensions may retain their popularity as a flexible means of decumulation—albeit unprotected from investment, inflation and longevity risks.

5.3.3 Financial Advice

Financial advisors can also leverage our research findings in guiding retiree decision-making. Chapter 4 showed that within the group of retirees favouring level or stepped drawdown amounts, 50% generate modest income streams of less than $5,800 annually. Presumably, the appeal of phased withdrawal accounts for these retirees lies in some combination of: investment freedom; bequest potential; precautionary savings; or other reasons. Regardless, using a phased withdrawal account to generate level or stepped income streams exposes the retiree to the risk of exhausting their account balance during retirement—due to favourable longevity experience or negative investment returns later in retirement.

Notably, in the absence of stepped annuities in the market, retirees can create an identical income stream arrangement through the purchase of two annuities: one guaranteed lifetime annuity at the 'stepped down' level, and a term annuity commencing immediately to generate higher income in earlier years. Alternatively, purchasing a term annuity and a deferred annuity, with different guaranteed levels of income, generates the same effect.

Moreover, the intention behind CIPRs is to design products that balance these needs for income, flexibility and risk management. Once superannuation funds begin offering CIPRs, financial advisors should assist members in planning the appropriate mixture of income streams and precautionary savings to adopt within these products.

5.4 Future Work

5.4.1 Extensions on Available Data

While this paper investigated the available data from APRA-regulated superannuation funds, a similar panel dataset produced by the Australian Tax Office covers a large sample of self-managed superannuation fund (SMSF) members. Applying the methodology from this paper to the SMSF data would allow researchers to comment on whether the findings from this project generalise to SMSF members.

To extend our methodology, future work can attempt to fit mixture models to the clustered data. Briefly, a mixture model would allow each cluster to have its own parametrisation of the regressors. For the group of individuals drawing 10% annually, for example, the regression equation would collapse down to a constant term with a value of 10%. In the level drawdown rate and unallocated cluster groups, however, the results would be nontrivial. Alternatively, researchers could obtain similar results by fitting, in turn, panel regression models to subsets of data for each of the observed clusters. Comparing the regression results across different clusters could then determine how the available regressors influence drawdown rates within a specific cluster.

As mentioned in Chapter 3, we removed from the dataset those individuals who died while under observation. Analysing drawdown behaviours in years immediately preceding death, and comparing the results against that of the surviving retirees in our sample, may reveal whether proximity to death significantly influences drawdowns.

Similarly, we calculated investment returns throughout retirement to construct the risk appetite metric as a time-invariant regressor. However, studying the evolution of investment returns within individual accounts throughout retirement may generate insights into how risk preferences change during retirement.

5.4.2 Remaining Gaps

A remaining gap is to determine how characteristics such as couple status, health and other wealth, which were not present in this study, influence financial decision-making in retirement. As superannuation funds are unlikely to collect or retain these variables, a more feasible approach would involve analysing data from Centrelink—the Australian government's welfare distribution service. Although drawdowns from account-based pensions may not be directly visible to the government, Centrelink may be able to combine information on the level of superannuation assets with the variables of interest, as well as Age Pension entitlements. A panel dataset of this nature could draw the link between the research in this paper on second-pillar behaviours with existing work on social security benefits in retirement (see e.g. Asher et al., 2017).

Furthermore, our results show that behavioural economics still has an important role to play in describing drawdown behaviours. Most owners of account-based pensions follow simple drawdown rules, and most members do not fully utilise the flexibility available in the product. Drawdown amounts rarely change in an inexplicable fashion, and the majority of members do not make adhoc drawdowns. These interesting behaviours motivate further explanation, and it is plausible that no selection of collected regressors could adequately predict adherence to a particular behavioural group. Instead, we support progressing the literature which empirically tests popular behavioural hypotheses, such as the impact of default options on financial decisions in retirement (see e.g. Bateman et al., 2017). In addition, collecting survey data on individuals who make adhoc drawdowns—why, when and how much—would provide a valuable contribution in explaining the large variability observed in the adhoc drawdown behaviours.

CHAPTER 6

CONCLUSION

6.1 Addressing Research Aim, Questions and Hypotheses

This paper began with the following aim:

Identify and explain drawdown behaviours in phased withdrawal products

Successfully achieving this aim is important for two reasons. First, it progresses the academic literature on drawdown behaviours within phased withdrawal accounts, which until now had relied primarily on theoretical studies into optimal behaviours, and lacked feedback from empirical studies. Second, it provides timely insights into appropriate policy decisions, retirement income product design, and financial advice, during a transitional period for Australia's retirement system.

By fulfilling this aim, we can now answer the three research questions posed initially.

1. *What drawdown behaviours are observed in account-based pensions?*
   The two most popular behaviours identified were: the default option of closely following the minimum drawdown rates; and drawing a level dollar amount over time, sometimes subject to downward revisions. Two other behaviours were present, although observed much less frequently. One of these involved drawing at a rate of 10% per year, which for some younger members is the upper bound on allowable drawdowns from TRIPs, and for other retirees may simply be an attractive heuristic. The second of the smaller behavioural groups was characterised by a complete drawdown of account balance while under observation.
2. *Are statistical models effective at predicting drawdown rates and behaviours?*
   Both panel data models and categorical regression models can provide insights into how member characteristics influence observed behaviours. However, there remains a large

portion of the total variation in observed drawdown rates and behaviours that could not be explained by the available characteristics.

3. *Which income products and policy design recommendations would suit the identified groups of retirees?*

   Retirees with a preference for level income streams, whether constant throughout retirement or subject to downwards revisions, could benefit from partial annuitisation, protecting them against the risk of outliving the assets supporting their income stream. By retaining a portion of their assets in a more liquid cash or investment account, they would still have money available to bequeath in case of early death, or alternatively as a source of wealth for adhoc withdrawals.

   Policymakers must recognise that the minimum drawdown rates strongly guide drawdown behaviours. Regular review of these rates is of paramount importance, as is caution when considering changes to the rates.

Furthermore, our methodology investigated several hypotheses, originating from the existing body of research. Here, we comment on these hypotheses with respect to the results found.

Annual Drawdown Rates

1. *Older individuals draw down less in excess of the minimum rates, compared to younger retirees*

   The behaviour appears to be different for different age ranges. Between the ages of 65 and 74, ageing decreases the excess regular drawdown rate, conditional on drawing in excess of the minimum rates. Beyond age 74, drawdowns in excess of the minima tend to occur at higher rates for older members. Note, however, that ageing concurrently decreases the likelihood of drawing above the minimum rates in the first place.

2. *Individuals with larger account balances draw less in excess of the minimum rates, compared to retirees with smaller account balances*

   Yes. Larger account balances give rise to smaller regular drawdown rates, and also smaller excess drawdown rates when drawdown is above the minimum. However, when members make adhoc drawdowns, those with larger account balances tend to have higher adhoc drawdown rates.

3. *Females draw more slowly through their account balances than males, after controlling for factors such as account balances*

   Yes.

4. *In financial years following the GFC, drawdowns in excess of the minimum rates decreased*

   No. In the financial years following the GFC, drawdown in excess of the minimum rates became more likely. Additionally, drawdown rates overall have tended to be higher since the GFC.

5. *In financial years following the GFC, the temporarily lower (concessional) minimum drawdown rates encouraged many retirees who had been drawing at the previous minimum rates to reduce their drawdowns to the concessional levels*

   Many, but not a majority. Just over one-third of retirees who used the minimum drawdown rates as a guide reduced their drawdowns to the concessional minima when they applied in financial years 2009–13. The remainder either preferred some combination of the concessional and non-concessional minima, or began to draw at much higher rates directly following the GFC.

Behavioural Groups in the Drawdown Series

1. *A substantial portion of retirees will draw consistently at minimum rates*
   Yes. Almost 50% of the observed members made close reference to the minimum drawdown rates. A quarter of these individuals, however, were disrupted around the time of the GFC, and were not able to recover.
2. *A group will attempt to draw at a constant rate, for example 7% per year*
   Few members drew at a constant rate over time, after accounting for drawdowns at the minimum rates. Within the 4% of our sample who drew regularly at a rate of 10%, 88% were in TRIPs, where this is the maximum allowable drawdown rate.
3. *Some will draw a constant nominal—not rising with inflation—dollar amount throughout retirement*
   Yes. Drawing level amounts was the second-most common drawdown behaviour.
4. *A group will draw a constant real—rising with inflation—dollar amount*
   No. We did not find evidence for this behaviour within account-based pensions.
5. *Some retirees will spend more than the minimum rates initially, but over time reduce drawdowns*
   Yes. Of those preferring to draw level amounts over time, just over one third revised down the level of their income stream during observation. Furthermore, for those not drawing level income streams, drawing at the minimum rates became more likely as members aged—even after controlling for the effect of rising minimum drawdown rates. Many plausible explanations exist for this behaviour, including: reduced consumption at older ages; desire to preserve capital for older ages; and bequest motives.

6.2 Summary

Chapter 1 opened this paper by contextualising the decumulation phase of retirement. In recent decades, demographic trends have driven larger employers to shift the responsibility of financial risk management in retirement to the former employees themselves. Several factors compound the difficulty inherent in making suitable choices on the threshold of retirement—including myopic thinking, financial illiteracy and susceptibility to cognitive biases. Due to the widespread use of phased withdrawal accounts, the study of behaviours within these products plays a key role in understanding the decumulation of assets in retirement. In addition, the Australian government has begun to focus on increasing levels of annuitisation by relaxing regulations and promoting CIPRs—which encourage a longevity-protected income component. Consequently, policymakers, financial advisors, and retirement income product designers can benefit from deeper insights into how retirees behave within account-based pensions. Collectively, these contextual factors motivated this research.

Next, Chapter 2 examined a broad literature on the decumulation phase of retirement. In particular, the studies in this field have investigated how retirees should, can, and do, draw down their accumulated wealth in retirement. Crucially, this chapter identified a gap in the literature. Despite several papers exploring suggested behaviours in phased withdrawal products, there has been a lack of adequate statistical analysis of the empirical drawdown rates and behaviours within these products. This analysis is necessary to determine the extent to which individuals utilise the heuristics suggested by the literature, and to identify any behaviours not yet considered. The identification of novel behaviours in retirement extends the theoretical literature by motivating further study into how retiree preferences drive the uncovered behaviours.

Chapter 3 detailed a methodology to fulfil the research aims and fill this literature gap, and Chapter 4 presented the results of applying this methodology to the available industry-level data from Australian superannuation funds. First, panel regression models relate drawdown rates to member characteristics. These models indicate the direction, magnitude and statistical significance of the effects of the regressors on several dependent variables. Second, a cluster analysis allocates members into distinct behavioural groups—characterised by their observed drawdowns over time. Third, a categorical regression model finds the statistical relationships between member characteristics and the likelihood of belonging to the identified behavioural groups. Additionally, investigations into the distribution of regular and adhoc drawdowns within particular groups reveal further insights into drawdown behaviours.

Finally, Chapter 5 discussed the results with respect to filling the identified gap in the literature, as well as the immediate social impact of these findings. Broadly, older retirees are more likely to draw at the minimum rates, and more likely to make adhoc drawdowns. They draw at slower rates when making regular drawdowns, but put more strain on their account balances when making adhoc drawdowns. Retirees with higher account balances tend to have slower regular drawdown rates, but draw through balances faster when making adhoc drawdowns. When facing higher minimum drawdown requirements, members are more likely to draw at the minima and less likely to make adhoc drawdowns. Their regular drawdown rates are higher, and they put less strain on account balances via adhoc drawdowns. In general, males draw down their account balances at faster rates than females, as do individuals with higher risk appetites and those who retired older.

Within the literature on drawdown behaviours, a valuable contribution from this work is finding that the large majority of our sample used two simple rules in retirement: following the minimum drawdown rates; or drawing level dollar amounts. Members who referenced the minima were more likely to be female, have larger account balances, a lower risk appetite, and have retired later. By contrast, retirees who drew constant amounts were more likely to be male, have smaller balances, a higher risk appetite, and have retired younger. These differences, while statistically significant, were relatively small in magnitude. However, members with older accounts were noticeably more likely to draw at the minimum than members who had opened their accounts since the latest drawdown rules came into effect. Additionally, two smaller behavioural groups exist in the sample: those who drew 10% annually; and those drawing down their entire account balance while under observation.

These findings have implications for policymakers, retirement income product designers, and financial advisors. On the policy side, it is clear that the magnetism of the minimum drawdown rates—or their use as the default option by superannuation funds—draws a large proportion of retirees to use them as guides. As a result, the government must continue to regularly review these minima, and realise the widespread impact of changing them.

For the design of more advanced retirement income products, it is clear that stepped annuities could play an important role in the market, as a large group of retirees construct their own equivalents within account-based pensions already. Furthermore, super funds creating CIPRs should cautiously avoid setting defaults for determining the proportion of assets which will support income streams versus an allowance for adhoc withdrawals. Most individuals do not appear to make adhoc drawdowns at all—while amongst those who do, there is huge variability in the proportion of assets withdrawn ad hoc versus regularly.

Finally, many retirees show a clear preference for drawing level income streams from their accounts, but are missing out on the potential longevity insurance provided by partial annuitisation of their superannuation wealth. These individuals in particular could benefit from financial advice directing them to allocate a portion of their accumulated superannuation assets into an income stream—either level, or level with a step down later in retirement.

APPENDIX A

GLOSSARY

Term	Definition
DB	Defined-Benefit
DC	Defined-Contribution
SISR	*Superannuation Industry (Supervision) Regulations* 1994
GFC	Global Financial Crisis
Legacy Account	Accounts opened prior to 20 September 2007
CIPR	Comprehensive Income Product for Retirement
OLS	Ordinary Least Squares
PC	Pooled Cross-sectional
FE	Fixed Effects
RE	Random Effects
HT	Hausman-Taylor
CRE	Correlated Random Effects
TRIP	Transition to Retirement Income Product
AME	Average Marginal Effect
SMSF	Self-managed super fund

BIBLIOGRAPHY

Aegon (2016). The Golden Age of Retirement – does rising pensioner wealth mask future problems? Technical report, Aegon. https://www.aegon.co.uk/content/dam/ukpaw/documents/retirement-income-report.pdf.

Asher, A. (2012). The lifetime harvesting plan: smoothed annuities with sharing of mortality, and averaging of investment, risks. In *Financial Services Forum*, Melbourne. Institute of Actuaries of Australia.

Asher, A. (2015). Retirement Incomes – Issues and Next Steps. *Australian Journal of Actuarial Practice*, 3:93–97.

Asher, A., Meyricke, R., Thorp, S., and Wu, S. (2017). Age pensioner decumulation: Responses to incentives, uncertainty and family need. *Australian Journal of Management*, 42(4):583–607.

Australian Government Australian Bureau of Statistics (2017). 6401.0 - consumer price index, australia, sep 2017. http://www.abs.gov.au/AUSSTATS/abs@.nsf/DetailsPage/6401.0Sep%202017?OpenDocument.

Australian Government The Treasury (2016a). Development of the framework for comprehensive income products for retirement. https://consult.treasury.gov.au/retirement-income-policy-division/comprehensive-income-products-for-retirement/supporting_documents/CIPRs_Discussion_Paper_1702.pdf.

Australian Government The Treasury (2016b). Retirement Income Streams Review. https://www.treasury.gov.au/~/media/Treasury/PublicationsandMedia/Publications/2017/RetirementIncomeStreamsReview/Downloads/PDF/RISR_May_2016.ashx.

Bateman, H., Chomik, R., and Piggott, J. (2016). Retirement Income. In O'Loughlin, K., Browning, C., and Kendig, H., editors, *Ageing in Australia: Challenges and Opportunities*. Springer New York.

Bateman, H., Eckert, C., Iskhakov, F., Louviere, J., Satchell, S., and Thorp, S. (2017). Default and Naive Diversification Heuristics in Annuity Choice. *Australian Journal of Management*, 42(1):32–57.

Bateman, H. and Piggott, J. (2011). Too Much Risk to Insure? The Australian (non-) Market for Annuities. In Mitchell, O. S. and Piggott, J., editors, *Securing Lifelong Retirement Income: Global Annuity Markets and Policy*. Oxford University Press, Oxford.

Bateman, H. and Thorp, S. (2008). Choices and constraints over retirement income streams: Comparing rules and regulations. *Economic Record*, 84(s1).

Benartzi, S., Previtero, A., and Thaler, R. H. (2011). Annuitization puzzles. *The Journal of Economic Perspectives*, 25(4):143–164.

Beshears, J., Choi, J. J., Laibson, D., Madrian, B. C., and Zeldes, S. P. (2014). What makes annuitization more appealing? *Journal of Public Economics*, 116:2–16.

Blake, D., Cairns, A. J. G., and Dowd, K. (2003). Pensionmetrics 2: Stochastic pension plan design during the distribution phase. *Insurance: Mathematics and Economics*, 33(1):29–47.

Brancati, C. U., Beach, B., Franklin, B., and Jones, M. (2015). Understanding retirement journeys: Expectations vs reality. (November). `http://www.ilcuk.org.uk/images/uploads/publication-pdfs/Understanding_Retirement_Journeys.pdf`.

Broadbent, J., Palumbo, M., and Woodman, E. (2006). The shift from defined benefit to defined contribution pension plans – implications for asset allocation and risk management. Technical report, Committee on the Global Financial System (CGFS), Bank of International Settlements (BIS).

Brown, J. R. (2001). Private pensions, mortality risk, and the decision to annuitize. *Journal of Public Economics*, 82(1):29–62.

Brown, J. R. (2009). Understanding the Role of Annuities in Retirement Planning. In Lusardi, A., editor, *Overcoming the Saving Slump: How to Increase the Effectiveness of Financial Education and Saving Programs*. University of Chicago Press, Chicago, IL.

Brown, J. R., Kling, J. R., Mullainathan, S., and Wrobel, M. V. (2008). Why don't people insure late-life consumption? A framing explanation of the under-annuitization puzzle. *American Economic Review*, 98(2):304–309.

Challenger Limited (2016). Annual Report 2016. `http://www.challenger.com.au/group/part_1_-_annual_report_2016.pdf`.

Challenger Limited (2017a). Interim Financial Report 2017. `http://www.challenger.com.au/group/Challenger_Interim_Financial_Report_2017.pdf`.

Challenger Limited (2017b). Rates - challenger. `https://www.challenger.com.au/products/rates.asp`.

Chamberlain, G. (1982). Multivariate regression models for panel data. *Journal of econometrics*, 18(1):5–46.

Chi, J. T., Chi, E. C., and Baraniuk, R. G. (2016). k-POD: A Method for k-Means Clustering of Missing Data. *The American Statistician*, 70(1):91–99.

Clare, R. (2013). Changes to regulatory settings for financial products dealing with longevity. Technical report, The Association of Superannuation Funds of Australia, Sydney.

Coile, C. and Milligan, K. (2009). How household portfolios evolve after retirement: The effect of aging and health shocks. *Review of Income and Wealth*, 55(2):226–248.

Commonwealth of Australia (2010). Final report. `https://taxreview.treasury.gov.au/content/FinalReport.aspx?doc=html%2FPublications%2FPapers%2FFinal_Report_Part_1%2Findex.htm`.

De Nardi, M., French, E., Jones, J. B., Journal, S., February, N., and Jones, J. B. (2015). Why Do the Elderly Save? The Role of Medical Expenses. *Journal of Political Economy*, 118(1):39–75.

Denuit, M., Haberman, S., and Renshaw, A. (2011). Longevity-indexed life annuities. *North American Actuarial Journal*, 15(1):97–111.

Denuit, M., Haberman, S., and Renshaw, A. (2015). Longevity-contingent deferred life annuities. *Journal of Pension Economics and Finance*, 14(03):315–327.

Donnelly, C. (2015). Actuarial Fairness and Solidarity in Pooled Annuity Funds. *ASTIN Bulletin*, 45(1):49–74.

Donnelly, C., Guillén, M., and Nielsen, J. P. (2014). Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, 56(1):14–27.

Emms, P. (2010). Relative Choice Models for Income Drawdown in a Defined Contribution Pension Scheme. *North American Actuarial Journal*, 14(2):176–197.

Fong, J. H. Y., Mitchell, O. S., and Koh, B. S. K. (2010). Longevity Risk Management in Singapore's National Pension System. *The Journal of Risk and Insurance*, 78(4):961–981.

Ganegoda, A. and Evans, J. (2017). The australian retirement lottery: A system failure. *Australian Journal of Management*, 42(1):3–31.

Garcia, R., Renault, E., and Semenov, A. (2006). Disentangling risk aversion and intertemporal substitution through a reference level. *Finance Research Letters*, 3(3):181–193.

Goda, G. S. and Manchester, C. F. (2013). Incorporating employee heterogeneity into default rules for retirement plan selection. *Journal of Human Resources*, 48(1):198–235.

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

Horneff, W. J., Maurer, R. H., Mitchell, O. S., and Dus, I. (2008). Following the rules: Integrating asset allocation and annuitization in retirement portfolios. *Insurance: Mathematics and Economics*, 42(1):396–408.

Hulley, H., Mckibbin, R., Pedersen, A., and Thorp, S. (2013). Means-tested public pensions, portfolio choice and decumulation in retirement. *Economic Record*, 89(284):31–51.

Institute of Actuaries of Australia (2014). Review of retirement income stream regulation. `http://www.actuaries.asn.au/Library/Submissions/Superannuation/2014/140922TreasuryRetirementIncomeStream.pdf`.

Iskhakov, F., Thorp, S., and Bateman, H. (2015). Optimal Annuity Purchases for Australian Retirees. *Economic Record*, 91(293):139–154.

James, E. and Song, X. (2001). Annuities markets around the world: money's worth and risk intermediation. `http://www.cerp.carloalberto.org/wp-content/uploads/2008/12/wp_16.pdf?d2641f`.

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). *An Introduction to Statistical Learning with Applications in R*. Springer.

Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5):1449–1475.

Lockwood, L. M. (2012). Bequest motives and the annuity puzzle. *Review of Economic Dynamics*, 15(2):226–243.

Lusardi, A. and Mitchell, O. S. (2007). Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics*, 54(1):205–224.

MacDonald, B.-J., Jones, B., Morrison, R. J., Brown, R. L., and Hardy, M. (2013). Research and Reality: A Literature Review on Drawing Down Retirement Financial Savings. *North American Actuarial Journal*, 17(3):181–215.

Mercer (2016). Melbourne Mercer Global Pension Index. Technical report. `http://www.globalpensionindex.com/wp-content/uploads/MMGPI2016-Report.pdf`.

Milevsky, M. (2013). *Life annuities: an optimal product for retirement income*. CFA Research Foundation.

Mitchell, O. S., Poterba, J. M., and Warshawsky, M. J. (1999). New Evidence on the Money's Worth of Individual Annuities. *American Economic Review*, 89(5):1299–1319.

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

Murtagh, F. (1983). A survey of recent advances in hierarchical clustering algorithms. *The Computer Journal*, 26(4):354–359.

Murtagh, F. and Legendre, P. (2014). Ward's hierarchical agglomerative clustering method: which algorithms implement ward's criterion? *Journal of Classification*, 31(3):274–295.

Piggott, J. and Detzel, B. (2005). The Simple Analytics of a Pooled Annuity Fund. *The Journal of Risk and Insurance*, 72(3):497–520.

Plan For Life (2016). Report on the Australian retirement income market. Technical report, Plan For Life.

Poterba, J., Venti, S., and Wise, D. A. (2013). The drawdown of personal retirement assets: husbanding or squandering? *NBER Working Paper*. `https://economics.mit.edu/files/6318`.

Qiao, C. and Sherris, M. (2013). Managing systematic mortality risk with group self-pooling and annuitization schemes. *Journal of Risk and Insurance*, 80(4):949–974.

Richter, A. and Weber, F. (2011). Mortality-Indexed Annuities: Managing Longevity Risk via Product Design. *North American Actuarial Journal*, 15(2):212–236.

Sneddon, T., Reeson, A., Zhu, Z., Stephenson, A., Hobman, E. V., and Toscas, P. (2016). Superannuation drawdown behaviour. *The Finsia Journal of Applied Finance*, (2):42–53.

Spicer, A., Stavrunova, O., and Thorp, S. (2016). How Portfolios Evolve after Retirement: Evidence from Australia. *Economic Record*, 92(297):241–267.

Stamos, M. Z. (2008). Optimal consumption and portfolio choice for pooled annuity funds. *Insurance: Mathematics and Economics*, 43(1):56–68.

Stringer, R. (2011). Addressing the Market Failure in Post-Retirement Products – Legal Issues. Presentation, 23rd annual conference held by the Law Council Superannuation Committee: Super - A Paradise Lost?

Swamy, P. and Arora, S. S. (1972). The exact finite sample properties of the estimators of coefficients in the error components regression models. *Econometrica: Journal of the Econometric Society*, pages 261–275.

Van Ooijen, R., Alessie, R., and Kalwij, A. (2015). Saving behavior and portfolio choice after retirement. *De Economist*, 163(3):353–404.

Ward Jr, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244.

Wooldridge, J. M. (2010a). Correlated random effects models with unbalanced panels. *Manuscript (version May 2010) Michigan State University.*

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

Wooldridge, J. M. (2012). *Introductory Econometrics a Modern Approach.* Cengage Learning.

Wu, S., Bateman, H., and Stevens, R. (2016). Optimal Portfolio Choice with Health-contingent Income Products: The Value of Life Care Annuities. ARC Centre of Excellence in Population Ageing Research Working Paper 2016/17.

Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economics Studies*, 32(2):137–150.