Modelling optimal decisions for financial planning in retirement using stochastic control theory

Johan G. Andréasson

School of Mathematical and Physical Sciences

University of Technology Sydney

Thesis submitted for the degree of

Doctor of Philosophy in Mathematical Science

Sydney 2017

Declaration of authorship

I certify that the work in this thesis has not previously been submitted for a degree nor has it been submitted as part of requirements for a degree except as fully acknowledged within the text.

I also certify that the thesis has been written by me. Any help that I have received in my research work and the preparation of the thesis itself has been acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

Production Note:
Signature removed prior to publication.

Johan G. Andréasson
February 1, 2018

Acknowledgement

First and foremost I would like to express my sincere gratitude to my academic supervisor, Prof. Alex Novikov, who offered me a unique opportunity to study at UTS, as well as my industry supervisor, Prof. Pavel V. Shevchenko, who has provided tremendous guidance throughout. Every aspect of this thesis has benefited from their supervision. They have kept me on a straight path to finish my thesis, helped me to avoid publication pitfalls, and provided feedback when progress stagnated.

I gratefully acknowledge the funding sources that made my research possible. Special thanks go to the Australian Technology Network's Industry Doctoral Training Centre (IDTC), which provided me with the opportunity to bridge an academic PhD with an industry placed research project, including its past directors Dr. Matt Brown and the UTS IDTC Node Leader Associate Prof. Yakov Zinder. I thank the IDTC industry partner CSIRO for partly funding this research, identifying the research problem, and supplying me with resources to carry out the research.

Finally, I would like to thank Dr. Sachi Purcal, Dr. Xiaolin Luo and A/Prof. Anthony Asher for discussions on material related to Chapters 3 and 4, Dr. Nicolas Langrené for discussion related to Chapter 5 and A/Prof. Juri Hinz for feedback on dynamic programming.

Abstract

In this thesis, we develop an expected utility model for retirement behaviour in the decumulation phase of Australian retirees with sequential family status subject to consumption, housing, investment, bequest, and a government-provided means-tested Age Pension. We account for mortality risk and risky investment assets, and we introduce a “health” proxy to capture the decreasing level of consumption for older retirees. The model is calibrated using the maximum likelihood method with empirical data on consumption and housing from the Australian Bureau of Statistic’s 2009-2010 ‘Household Expenditure Survey’ and ‘Survey of Income and Housing’. The calibrated model fits the characteristics of the data well to explain the behaviour of Australian retirees, and is then used to examine the optimal decisions given recent Age Pension policies and different family settings. Specifically, we examine optimal decisions for housing at retirement, and the optimal consumption and risky asset allocation depending on age and wealth for the Age Pension policies 2015-2017.

As the piecewise linearity in the Age Pension function requires the stochastic control problem to be solved numerically, we utilise the Least Squares Monte Carlo method to extend the problem with additional dimensions and control variables. This method is difficult to use with utility functions, as it can lead to a bad fit or bias from transforming variables. We suggest methods to account for this bias, and show that the Least Squares Monte Carlo is then accurate when applied to expected utility stochastic control problems. We then extend the optimal decisions to include annuitisation, as well as the option to scale housing in retirement or to access home equity through a reverse mortgage, and examine optimal decisions with respect to the Age Pension in retirement.

Table of contents

List of tables

List of figures

Chapter 1

Introduction

Countries across the globe are experiencing increased longevity, and providing a public pension as the population ages remains challenging. In general, two types of pension plans exist for retirement: defined benefit and defined contribution. In a defined benefit pension plan, the pension payments are predetermined by a formula based on the retiree's earnings history. Defined contribution, on the other hand, is based on (mandatory) contributions from the retiree's salary where the pension assets can generate investment returns for the retiree. Defined benefit is known to have many problems, where the central issue is the difficulty in estimating the present value of the future liability of the pension plan (Rothman, 2012; Poterba et al., 2007). The present value depends on a number of variables, such as the length of working career, salary growth and mortality, where the risks associated with each variable are difficult to estimate. It can be greatly underestimated when longevity increases, which has been the case for the last century. As more nations realise the problems that come with defined benefit pensions and shift their policies (completely or partially) towards defined contribution, a new branch in the life cycle model research field is created which has been the focus of financial and actuarial publications in recent years. This thesis aims to develop a mathematical model for the decumulation phase (retirement) of a life cycle model, while accounting for practical matters such as a means-tested Age Pension and regulatory requirements. Such a model is characterised by a multi-dimensional stochastic control problem[1].

[1]Throughout the thesis the terms 'stochastic control problem', 'optimal control problem' and 'dynamic programming problem' will be used interchangeably. They all refer to multi-period problems where control (decision) is dependent on a certain state, and the problem is stochastic

The focus is on a life cycle model that captures the behaviour in retirement specific to Australian retirees and that can be used in simulation on an individual retiree level for optimal retirement planning, as well as to measure the impact of policy changes and financial products.

1.1 Background

The shift from a defined benefit pension system to a defined contribution pension system transfers risk from the corporate sector (or state) to households, primarily via decisions to investment and withdrawal pension assets. Although defined benefit schemes remain available, most are closed for new members and have been replaced with defined contribution schemes. As a result, Australia has accumulated Superannuation[2] assets of $2.02 trillion dollars (ASFA, 2015), making it the fourth-largest pension fund pool in the world. For retirees, the main difference is that instead of being provided with a monthly benefit as they were previously, they now receive a lump sum at retirement and are responsible for managing this wealth throughout their lives. Retirees face multiple risks that are difficult to account for, including mortality, longevity and investment risks, as well as regulatory risks such as changes in policies and government-provided Age Pension entitlements. The long-term effects of this new pension system remain unknown.

Many other countries that offer social pensions incorporate some kind of means-test, such as the USA (Supplemental Security Income[3]) and Greece (Social Solidarity Benefit[4]). However, Australia provides a unique pension system as it does not have an earnings-related scheme. Moreover, it currently has the largest mandatory defined contribution scheme. Public pensions in other countries tend to be more generous and increase with earnings, whereas those in Australia reduce with means. In addition, the Australian system is still maturing; a few decades will pass before all retirees have been through the Superannuation system. With around 72% of the Australian population aged 65 or older being entitled to a full or partial Age

in nature.

[2]In Australia, the arrangement for people to accumulate funds for income in retirement is referred to as *Superannuation*.

[3]https://www.ssa.gov/ssi/

[4]www.oecd.org/greece/

Pension[5], this puts financial pressure on the government. Therefore, there is an urgent need for a model that not only captures the characteristics of Australian empirical data but also explains the behaviour of Australian retirees.

1.1.1 The Australian retirement system

Australia has adopted the three pillar system outlined by World Bank (2008), and relies on a defined contribution pension system that is based on the Superannuation Guarantee, private savings, and a government provided Age Pension. These three pillars are adopted from the The Superannuation Guarantee mandates that employers contribute a set percentage of the employee's gross earnings to a Superannuation fund, which accumulates and is invested until retirement. The current contribution rate is set to a minimum[6] of 9.5%, supporting both defined benefit and defined contribution plans, where contributions in addition to this often come with tax benefits. Private savings are comprised of these voluntary Superannuation contributions, but also include savings outside the Superannuation fund such as investment accounts, dwellings, and other assets. Finally, the Age Pension is a government managed safety net which provides the retiree with a means-tested Age Pension.

The means-test determines whether the retiree qualifies for full, partial, or no Age Pension once the entitlement age of 65.5 years is reached. In the means-test, income and assets are evaluated individually, and a certain taper rate reduces the maximum payments once the income or assets surpass set thresholds (which are subject to family status and home ownership). The exact amount of Age Pension a retiree qualifies for is determined by the smaller outcome of the income and asset-tests. Income from different sources are also treated differently; financial assets and new account-based income streams (after the 1st of January 2015) are expected to generate income based on a progressive deeming rate, while income streams such as labour and non-account-based annuity payments or existing account-based income streams are assessed based on their nominal value. Deemed income refers to the assumed returns from financial assets without reference to the actual returns on the

[5] In June 2014, 2,404,902 people received an Age Pension (Department of Social Services, 2014) out of a total population of 3,337,502 aged 65 or older (Australian Bureau of Statistics, 2014).

[6] Some industries have higher rates, e.g., universities pay a 14% Superannuation contribution for employment longer than 12 months.

assets held, and is calculated as a progressive rate of assets. Two different deeming rates may apply based on the value of the account, a lower rate for assets under the deeming threshold and a higher rate for assets exceeding the threshold, as shown in Table 1.1. In addition to the general structure of the Age Pension and means-tests, there are exceptions to these rules due to things like disabilities, illness in couple households, and exact year of retirement where the retiree is affected by significant changes to the system (transitional arrangements).

Table 1.1: Age Pension rates published by Centrelink as at June 2017 (*www.humanservices.gov.au/customer/services/centrelink/age-pension*).

	Single	Couple
Full Age Pension per annum	$22,721	$34,252
Income-Test		
Threshold	$4,264	$7,592
Rate of Reduction	$0.5	$0.5
Asset-Test		
Threshold: Homeowners	$250,000	$450,000
Threshold: Non-homeowners	$375,000	$575,000
Rate of Reduction	$0.078	$0.078
Deeming Income		
Deeming Threshold	$49,200	$81,600
Deeming Rate below Threshold	1.75%	1.75%
Deeming Rate above Threshold	3.25%	3.25%

Certain account types for retirement savings have a minimum withdrawal rate once the owner is retired. The most popular type, the Allocated Pension account[7], is an account that has been purchased with Superannuation and generates an income stream throughout retirement. It has the benefit that investment returns are tax-free, as contributions are already taxed. As with other income stream based accounts, it is subject to regulatory minimum withdrawal rules. The retiree must therefore withdraw a certain percentage each year, which increases with age, to avoid penalties or additional taxes (Table 1.2). The purpose of minimum withdrawals is to exhaust the retiree's account around year 100, in order to avoid the retiree living off the Age Pension and bequeathing the assets rather than supporting the retirement.

[7]The terms Allocated Pension account, account-based pension, phased withdrawal products or income stream based accounts refer to the same account structure and are used interchangeably in the thesis.

Table 1.2: Minimum regulatory withdrawal rates for Allocated Pension accounts for the year 2017 and onwards (*www.ato.gov.au/rates/key-superannuation-rates-and-thresholds*).

Age	≤ 64	65-74	75-79	80-84	85-89	90-94	$95 \leq$
Min. drawdown	4%	5%	6%	7%	9%	11%	14%

1.1.2 Revisions and changes to the system

Since the Australian retirement system is relatively young, the long-term effects of this new pension system are not yet known. Changes in this system are expected to occur frequently as a result of fiscal decisions, and as the impact that policy changes have on a retiree's personal wealth (and the economy in general) becomes evident. Variables directly related to the means-test such as entitlement age, means-test thresholds, taper rates, and pension payments can all be adjusted to meet budget needs by the government. On a larger scale, regulatory changes may include whether the family home is included in the means-tested assets, the elimination of minimum withdrawal rules, changes in mandatory savings rates, or additional taxes on Superannuation savings. From a mathematical modelling perspective, this poses difficulties in terms of future model validity, as regulatory risk and policy changes can quickly make a model obsolete if it is not modified to account for the new rules.

The thresholds for the asset and income-test are adjusted yearly (Department of Social Services, 2016). The deeming rates are subject to change in relation to interest rates and stock market performance. The current rates are at a historical low. In 2008 the deeming rates were as high as 4%/6%, but in March 2013 they were set to 2.5%/4% due to decreasing interest rates, then in November 2013 to 2%/3.5% and to current levels of 1.75%/3.25% in March 2015. The Age Pension payment rates are indexed in accordance with the greater movement of the Consumer Price Index (CPI) and the 'Pensioner and Beneficiary Living Cost Index'[8] (PBLCI, and adjusted on the 20$^{\text{th}}$ of March and September each year). If the annual payment rates after CPI and PBLCI indexation are lower than 27.7% of the annualised 'Male Total Average Weekly Earnings' (MTAWE) index for singles, or lower than 41.8% for couples, then the annual payment rates are instead reset to 27.7% and 41.8% of

[8] The index is produced by the Australian Bureau of Statistics, and measures the price changes of out-of-pocket expenses for the subgroup 'age pensioner households and other government transfer recipient households'.

MTAWE respectively.

The research in this thesis has been dynamic due to the constant changes. As the period covered starts at 2010, where the most recent data for calibrating the model come from, up to the most recent 2017 change, there are some major changes the reader needs to be aware of.

2008 - In response to the global financial crisis of 2007–2008, the government provided pension drawdown relief by reducing the minimum withdrawal rates between 2008 and 2013. From the financial year 2008–2009 up to 2010–2011 the minimum withdrawal rates were decreased by 50%, and up to 2012–2013 decreased by 25% before returning to the default rates in 2013–2014 and onwards (Department of Social Services, 2016).

2015 - Before the 1st of January, income streams (such as Allocated Pension accounts) allowed for an 'income test deduction', which was decided on when the account was opened. This deduction effectively increased the threshold for a full Age Pension in the income-test, and remained constant each year. In addition, the means-tested income consisted of labour income and drawdowns from account-based pensions. After the 1st of January, the income-test deduction was no longer able to be used for new Allocated Pension accounts, and drawdown from such accounts was no longer considered income. Instead, the accounts were determined to generate income by a deeming rate that was applied to the assets. Already opened accounts, however, were 'grandfathered' and would continue to be assessed under the old rules (Department of Social Services, 2016).

2017 - The thresholds for the asset-test were 'rebalanced', so that the thresholds were increased but the taper rate doubled, as a means of controlling who is entitled to partial or full Age Pensions. The higher thresholds allow a retiree to hold more wealth while still receiving an Age Pension. The rebalancing will improve fairness and affordability of the pension system, where more than 90% of retirees are expected to be better off or not affected (Australian Government Department of Veterans' Affairs, 2016).

Future changes - The currently planned changes, in addition to adjusting Age Pension payments and thresholds accordingly, refers to an increase in the compulsory Superannuation contribution rate by 0.5% yearly starting 2021–2022, to reach 12%

in 2025–2026 (Australian Taxation Office, 2017). The entitlement age increased to 65.5 years on the 1st of July 2017, and will increase by six months every two years to reach 67 years in 2023. The tax exemption that is currently available for Allocated Pension accounts will be extended to other types of retirement products, such as deferred lifetime annuities and group self-annuities[9].

1.2 Literature review

While utilitarianism[10] stretches as far back as ancient times, the foundation for modern day utility theory was established by Neumann and Morgenstern (1944). Life cycle modelling (LCM) based on utility theory originates from Fisher (1930) and was later updated by Modigliani and Brumberg (1954) who observed that individuals make consumption decisions based on resources available at the current time, as well as over the course of their lifetimes. Essentially, LCM is a way to model an individual's consumption and savings preferences over the long term, with the overarching objective of maintaining one's current lifestyle. LCM reflects a large number of interrelated problems in a model that needs to be solved in order to find a solution and quickly becomes mathematically challenging. The key work for early models was laid out by Yaari (1964, 1965) who extended the LCM to uncertain lifetime and studied the optimal choice of life insurance and annuities. The original model in Fisher (1930) later led to the models of Samuelson (1969) and Merton (1969, 1971) who studied the problem in relation to optimal portfolio allocation, where Richard (1975) derived a closed-form solution for the consumption and investment decisions of an uncertain lifetime in a continuous time model. The solution was extended by Karatzas et al. (1986) to include general utility functions.

There is nowadays a vast literature that focuses on LCM or is indirectly related to the optimal consumption problem, such as portfolio allocation, mortality modelling, annuitisation, and other financial products to mention a few. While it is not feasible to delve into every minor extension of LCM, it is of interest to consider different solution techniques of the above mentioned problems as their solution remains chal-

[9]Participants contribute funds to a pool invested in financial assets, where payments are made to surviving members. This pools mortality risk and provides longevity protection.

[10]An ethical theory where the aim of action should be the largest possible balance of pleasure over pain, hence the best action is the one that maximises utility.

lenging in practical applications. Mainly, models can be either discrete or continuous in time, and a solution can be found analytically or numerically. Analytical solutions are available only in special circumstances, depending on the stochastic nature, dimensionality, and real-world constraint complexity. Analytical techniques used to study the retirement life cycle belong to two main categories: duality (martingale) based methods, or dynamic programming. In the duality approach, a constrained problem is transformed into an unconstrained dual 'shadow price'[11] (for applications in retirement, see for example Lim et al. (2008); Shin et al. (2007)). The dynamic programming method is what was originally used in Merton (1969), and is more common with discrete time solutions (Yao et al. (2014); Cocco et al. (2005); Bateman et al. (2007) to name a few). The main idea is to rewrite the multi-period problem so that optimal control is solved for a single time period via the Bellman equation and then solved recursively (Bellman, 2003), so that the optimal control for the full period can be found. By knowing that optimal control actions for certain states will be taken in the future, the optimal control can be found for present time. In the case of continuous-time problems, the analogous equation is the partial differential equation (PDE) often referred to as the Hamilton-Jacobi-Bellman equation (see for example Huang et al. (2012); Shin (2012)) which relies on stochastic calculus for an analytical solution. If the optimal strategy is non-linear, or piece-wise linear as in the case of modelling the retirement phase accounting for the Australian Age Pension, the solution must be computed numerically.

Regardless of solution type, it comes with drawbacks - mainly that it suffers from the curse of dimensionality[12]. PDE methods are only practical up to two dimensions and limit the dynamics of stochastic variables if an analytical solution is required. Even if numerical integration solutions, such as deterministic quadrature based ones, can in theory handle many dimensions, the computational requirements quickly become infeasible. These methods typically discretise the state and control on a grid, and calculate the expected value function for the next period using quadrature and interpolation between points (Luo and Shevchenko (2014); Hulley

[11]The instantaneous change (marginal utility) in the objective function by relaxing a constraint.

[12]The curse of dimensionality refers to complexities which arise in high-dimension space that is not present in low-dimension spaces. In stochastic control problems such complexity appears when the number of state variables increase, then the number of state-space combinations in the problem increases exponentially.

et al. (2013); Cocco and Gomes (2012) to name a few). As the number of states and stochastic factors increase, the calculations become unmanageable even for modern high-performance computing. There are methods that try to decrease the computational effort while keeping the same accuracy. One such method is the 'Endogenous Gridpoint Method', which re-writes the total amount of current states into a single future state variable at the end of the time period (Carroll, 2006), and extended in, e.g., Barillas and Fernández-Villaverde (2007) and Iskhakov (2015). The method can lead to closed form solutions of first-order conditions which speed up calculations significantly, but is dependent on the dynamics of the model and must be set up as a system of first-order conditions which can be solved with root-finding operations. Another method of interest is the Expectation-Maximisation-Control (EM-C) algorithm by Kou et al. (2016), which sequentially updates the control policies using Monte Carlo simulation by iterating a forward simulation and backward solution. The method can be applied to utility based stochastic control problems and is monotonic in improving the control for each iteration, but falls short in that it requires a parametric optimal control which is not suitable with respect to the Australian Age Pension.

A more promising research field is Approximate Dynamic Programming (ADP) methods, which have become increasingly popular. ADP methods, such as Least Squares Monte Carlo (Longstaff and Schwartz (2001); Tsitsiklis and Van Roy (2001); Kharroubi et al. (2014) to name a few), Convex Switching Systems (Hinz, 2014; Hinz and Yee, 2017) for optimal switching problems, or the Stochastic Mesh method (Broadie and Glasserman, 2004), can better deal with high dimensions in state, control, and disturbance space although at the expense of decreased accuracy. The central idea is to approximate the value function, and use this approximation to solve the optimal control. As the solution is an approximation only, it will always be sub-optimal, and it is difficult to measure exactly how sub-optimal the solution is. ADP methods are broadly classified into Value Function Iteration (VFI) or Policy Function Iteration (PFI) and will be further reviewed in Chapter 5.

1.2.1 Australia specific research

Australia is a unique case of means-testing, where there are two main approaches for modelling the Age Pension: expected utility maximisation via dynamic programming (Hulley et al., 2013; Bateman et al., 2007; Ding, 2014; Iskhakov et al., 2015) and overlapping generation (OLG) models based on dynamic general equilibrium (Kudrna and Woodland, 2011a,b, 2013; Tran and Woodland, 2014; Kudrna, 2014; Cho and Sane, 2013). Although both methods involve utility maximisation, they differ in terms of scope. While OLG models tend to analyse the effects on the economy in general by using a population of retirees, dynamic programming models focus on the retiree as an individual. With respect to Australia, the OLG models focus on how policy changes affect the welfare and economy-wide implications, even though household level decisions with respect to utility maximisation are part of the model. These micro-level decisions are very basic where state variables have few discrete steps, hence the accuracy of such solutions is rough and not very useful for deriving optimal behaviour for a retiree. Kudrna and Woodland (2011b) only allow for a few wealth profiles, while Tran and Woodland (2014) discretise the wealth state. As these models impose many assumptions and restrictions on the dynamics, which is required for the numerical solutions to be feasible (the means-test does not allow for analytical solutions), the optimal behaviour on the individual level is highly approximate. In addition, OLG models tend to have a very limited number of state variables and lack stochastic variables, again due to calculation limitations if analytical solutions cannot be found. However, owing to the overlapping component, OLG explicitly models various periods of the life cycle, thus generally including both the accumulation phase and the decumulation phase. As multiple generations are evaluated, such models are therefore useful to test the changes to policies upon impact, as they transition and eventually become long term effects. Kudrna and Woodland (2011a) and Tran and Woodland (2014) evaluate the effect of the means-tests and taper rate changes on the welfare outcome, while Kudrna and Woodland (2011b) examine the implication of the recent Australian pension reform extended with a housing market, which was further extended by Kudrna (2014).

Contrary to the OLG model, it is common to model decisions and behaviour in retirement on a micro-level using expected utility maximisation models based on dynamic programming techniques. Such methods are suitable to find the retirees' optimal decisions in the life cycle, as they allow for more complex dynamics and control variables than would be realistic to include in OLG models. Researchers have mainly focused on the economic effects of different means-tested policies and the impact on savings (Hubbard et al., 1995; Hurst and Ziliak, 2004; Neumark and Powers, 1998 to name a few), or how annuitisation is affected by the means-test (Iskhakov et al., 2015; Bütler et al., 2016). Typically, the kind of utility used is constant relative risk aversion (CRRA) or hyperbolic absolute risk aversion (HARA). HARA is preferred to CRRA in most studies because of the presence of a consumption floor, which implies a simplified preference for risk as the consumption floor needs to be maintained, although making the optimal decision dependent on wealth. With regard to CRRA, the retiree receives utility from absolute consumption, which does not fit the empirical data of consumption well for less wealthy households. Hulley et al. (2013) investigate the effect of the Age Pension on consumption and investments under CRRA utility, while Bateman et al. (2007) compare the effect of modelling optimal consumption and investment with HARA utility against CRRA, and argue that CRRA oversimplifies risk attitudes. Iskhakov et al. (2015) model the behaviour with a HARA utility maximisation approach in order to investigate how annuity purchases are affected by different preferences and scenarios.

As the Age Pension is piece-wise linear due to the means-test, Australian LCM needs to resort to numerical solutions. To our knowledge, it is only Ding (2014) who makes an attempt for a (semi) analytical solution, where the solution is analytical assuming the time periods for the different means-test phases (no Age Pension, partial due to the asset/income-test and full Age Pension) are correctly estimated. If not, the analytical solution needs to be iterated until it converges and therefore becomes semi-analytical. The model captures the behaviour of Australian retirees with regard to consumption, portfolio allocation, and housing using a HARA utility model.

There is no shortage of empirical research on Australian retirees based on various types of regression models (Higgins and Roberts, 2011; Olsberg and Winters, 2005; Asher et al., 2017; Lim-Applegate et al., 2006; Hulley et al., 2013). However, without a proper model that captures the behaviour of the retirees, the impact of changes to the pension system can only be evaluated after it has been implemented. This is known as the "Lucas critique", who argued that estimated parameters that are regarded as structural actually depends on the economic policy, making it useless for predicting effects from policy changes (Lucas, 1976). There are some attempts to calibrate models to Australian retirement behaviour, although most of them fall short on an individual level. Tran and Woodland (2014) and Kudrna and Woodland (2013, 2011b) calibrate their models to the Australian economy, but do not calibrate retiree preferences such as risk aversion. Ding (2014) calibrates the model with data from the Australian Bureau of Statistics (2011) using Mean-Square Error minimisation, although normalised over expected lifetime wealth which allows for larger errors in less wealthy households. Iskhakov et al. (2015) parameterise the model with variables extracted from Australian empirical data, but do not calibrate the model intrinsically.

A specific area of interest related to Age Pension modelling is the impact of policy changes. Bateman and Thorp (2008) evaluate how the changes of the drawdown rules introduced in 2007 affect the welfare of retirees. Ding (2014) examines how future projected costs are impacted by policy changes with regard to changes in the Superannuation Guarantee, asset-test treatment of housing, and indexation of Age Pension payments to price inflation. Cho and Sane (2013) evaluate policies related to including the family home in the means-test (although from a macroeconomic perspective), and Rothman (2012) utilises the RIMGROUP model to simulate forecasts on a range of different policy changes (changes in the Superannuation Guarantee, tax concessions, entitlement age, etc.) in order to estimate the impact on long-term government costs. The RIMGROUP model, which is a cohort projection model, plays a role in actual policy decisionmaking (Australian Department of Treasury, 2010). The model is developed by the Australian Treasury, thus specific information about the model is not publicly available. It is based on separate projections for each 'birth year gender decile cohort', which starts with labour force models and tracks/estimates changes in Superannuation accounts, savings accounts, tax liabilities, and so on. It utilises close to 100,000 records, which are divided into different subgroups in order to create projections from changes and trends in the population, but as it is population based it does not allow for micro-simulations.

1.2.2 Summary

In summary, the research gaps in previous research are as follows:

- Current models are limited with respect to stochastic factors and control variables, which tend to have a significant impact on optimal decisions. The lack of stochasticity in a utility function will greatly underestimate risks. The limitations are due to computational restrictions with numerical solutions, or constraints on dynamics and states for analytical solutions.

- Current research lacks realistic modelling on a micro-level, as well as the impact of policy changes on the retiree as an individual. This impact in conjunction with the means-test is a rather unexplored area in research where contributions would be greatly beneficial for many different decision-makers.

- Calibration is rather limited, or has been carried out in such a way that allows for large errors in less wealthy households. For a model to be useful in regard to the means-test, it is crucial that the subset of the dataset where the means-test binds has a better fit than the subset of people who do not receive any Age Pension.

- Age Pension policies are constantly changing, hence current research quickly becomes outdated (e.g., Hulley et al. (2013) used the 2006 policy, Bateman and Thorp (2008) the 2007 policy change). While some findings are less sensitive to the active policy, others are highly sensitive and need to be re-evaluated with the updated policy rules. The recent changes in the Age Pension policies have therefore not been examined in published research yet.

1.3 Objectives and scope

This research aims to model the retirement phase of Australian retirees with respect to the means-tested Age Pension in a more realistic framework than currently available. The thesis is built around an open-ended research question: '*Can a framework that captures the behaviour of Australian retirees with reasonable accuracy be defined, such that the model represents a realistic risk environment and is applicable on both* *a micro and macro scale?'*. The research will therefore focus on an individual retiree level that allows multiple stochastic factors and optimal control variables.

The aim can be divided further into sub-problems. As a start, a model suitable for Australian retirees needs to be defined. This model will then need to be extended to include risk variables and decisions that represent a realistic financial environment for the retiree. Finally, methods of solving the problem mathematically need to be evaluated and/or constructed, and the level of realism in the problem needs to be balanced with the computational time required to find a solution. In order to solve the general research aim and each particular sub-problem, the following is need:

1) to define a model that captures the characteristics of Australian retirees, and calibrate the model with empirical data to ensure its validity.

2) to investigate issues related to the Australian means-test, in order to propose suitable solution methods for the stochastic control problem owing to the piecewise linearity in the pension function.

3) to balance the model between complexity and computational efforts. Introducing additional states will lead to the curse of dimensionality, and introducing additional optimal control variables or stochastic factors will increase the computations needed during optimisation.

Due to the second point above, we are forced to utilise numerical methods to solve the life cycle problem. This, in turn, puts more emphasis on the third point, which is essentially the limiting factor for how realistic the model can be.

1.4 Limitations

The research will not model retirement from a macroeconomic perspective nor overlapping generations. While the model presented can actually be used in a macro sense by simulating a population and aggregating individual retirement behaviour to represent the macro effect, this is a positive byproduct and no attempts will be made to fit this to moments of macro variables nor explain any macro effects.

In addition, the research focuses on characteristics and effects related to the majority of the retirees. As described in Section 1.1.2, the Australian Age Pension is a highly dynamic system which continuously changes and where many exceptions and specific rules apply to sub-groups of entitled retirees. These minorities have either been omitted from any data sets or are assumed to behave like the majority of Australia's retirees. The model therefore may not allow their individual situation to be modelled, or alternatively may require a different parameterisation in order to do so.

1.5 Research significance

The thesis contributes to the fields of research of LCM, especially in the niche field of Australian retirement modelling, but also to adaptive dynamic programming in general. First, the model presented is the most realistic one to have been published on Australian Age Pension modelling to our best knowledge. It provides a better fit for empirical data, includes housing in the bequest, and allows a couple's family status to be sequential in order to account for change over time due to death of a spouse. It applies fewer restrictive assumptions on the dynamics of the means-test, and allows for more state, control and stochastic variables.

Second, this research is among the few to apply the Least-Squares Monte Carlo method to LCM, and to our best knowledge the only one who effectively avoids suffering from a re-transformation bias or bad regression fit when applied to utility functions (this will be explained further in Chapter 5). A suggested solution is presented to deal with common problems related to utility functions, which significantly improves the accuracy and allows us to extend the model with more complex dynamics.

Finally, on the practical side, the research expands the current work on policy changes and its impact on the behaviour of Australian retirees. Recent policy changes are evaluated on a retiree level rather than from a macro perspective. We examine the options of annuitising, up or downscaling one's home, or accessing home equity through reverse mortgages in retirement, which previously have either not been done with respect to Australia or done in a more limited framework.

While it is not realistic to specify a comprehensive framework for modelling the Australian Age Pension on both micro and macro scale, the suggested methods and model in the thesis will lend themselves to both perspectives. The model is expected to be used in:

- financial planning on an individual retiree level, where the model can be adapted to individual risk preferences.
- developing new retirement products by analysing how the behaviour of retirees changes, and whether it can create additional utility.
- large-scale simulation, as a component thereof, to estimate the effect of policy changes on both the individual retiree as well as government budgets.

1.6 Thesis structure

This introduction presented the background of the Australian Age Pension, an overview of the related literature, and how the thesis relates to the research field overall. Note that the literature review only provided the overlapping literature in relation to the thesis in general, to avoid repeating material which relates to more than one chapter. Literature reviews specific to each individual chapter are discussed in the introduction of the respective chapter. The thesis is based on four papers, where Andreasson et al. (2017) and Andréasson and Shevchenko (2017) relate to Chapters 3 and 4, Andreasson and Shevchenko (2017a) relates to Chapter 5 and Andreasson and Shevchenko (2017b) relates to Chapter 6.

Next, in Chapter 2 the required mathematical background is provided, in order to understand utility theory and the dynamic programming framework used to solve the stochastic control problem. Chapter 3 then presents the model in detail and structures it as a stochastic programming problem, but only provides a solution with arbitrary parameters. Chapter 4 focuses on the calibration of the model, and analyses the results once calibrated with recent policy changes. Chapter 5 lays out the foundation of the Least-Squares Monte Carlo algorithm used to solve a more flexible model, and suggests an approach to deal with problems related to utility functions. This chapter is crucial to be able to extend the model further. Chapter 6 then applies the Least-Squares Monte Carlo solution to extended models that conclude either annuitisation or additional control variables with respect to home equity. Finally, a concluding discussion is presented in Chapter 7 together with major findings and how the model can be applied in the finance industry.

There are four appendices located in the back. Appendix A explains the data aggregation process to create the samples used in the calibration. Appendices B and C state the techniques used to deal with re-transformation bias, and appendix D shows how to obtain the analytical solution used in examples in Chapter 5.

1.7 List of publications

Andréasson, J. G., Shevchenko, P. V., Novikov, A. (2017). Optimal consumption, investment and housing with means-tested public pension in retirement. *Insurance: Mathematics and Economics 75*, 32-47.

Andréasson, J. G., Shevchenko, P. V. (2017). Assessment of policy changes to means-tested Age Pension using the expected utility model: Implication for decisions in retirement. *Risks 5*(3).

Andréasson, J. G., Shevchenko, P. V. (2017). Bias-corrected Least-Squares Monte Carlo for utility based optimal stochastic control problems. Preprint available at SSRN: 2985828.

Andréasson, J. G., Shevchenko, P. V. (2017). Optimal annuitisation, housing decisions and means-tested public pension in retirement. Preprint available at SSRN: 2985830.

1.8 List of presentations

Andréasson, Johan G. A Utility Maximising Framework for Australian Age Pension, *Stochastic Methods in Quantitative Finance and Statistics Seminar Series*, Sydney, Australia, 2015.

Andréasson, Johan G. Optimal Decisions for Consumption, Investment and Housing in the Australian Retirement Decumulation Phase, *Quantitative Methods in Finance 2015 Conference*, Sydney, Australia, 2015.

Andréasson, Johan G. Optimal Decisions for Consumption, Investment and Housing in the Australian Retirement Decumulation Phase, *Stochastic Methods in Finance, Insurance and Statistics*, Shoal Bay, Australia, 2015.

Andréasson, Johan G. Optimal Decisions for Consumption, Investment and Housing in the Australian Retirement Decumulation Phase, *2016 ASTIN Colloquium*, Lisbon, Portugal, 2016.

Andréasson, Johan G. A Least-Squares Monte Carlo Method in Retirement Modelling, *Quantitative Methods in Finance 2016 Conference*, Sydney, Australia, 2016.

Andréasson, Johan G. A Least-Squares Monte Carlo Method in Retirement Modelling, *Stochastic Methods and Control Workshop 2017*, Trier, Germany, 2017.

Chapter 2

Mathematical background

In this chapter, we introduce some notation and describe various mathematical objects necessary for the formulation of the arguments used in this thesis. A complete framework for stochastic control problems is also defined, as the solution of the models introduced in the thesis relies on this framework. If the notation for expectation $\mathbb{E}[\cdot]$ is used, then it is assumed that the expectation exists.

2.1 Utility theory

Utility theory explains the behavior of individuals based on the premise that people can consistently rank order their choices depending upon their preferences. In order to state this mathematically, we start with the definition of convexity and concavity in functions, and a few preference relations.

Definition 1 (Convex functions)**.** *Let $f(x)$ be a real valued function defined on the interval $I = [a, b]$. The function f is said to be convex if for every $x_1, x_2 \in [a, b]$ and $0 \geq \lambda \geq 1$,*

$$f(\lambda x_1 + (1 - \lambda)x_2) \leq \lambda f(x_1) + (1 - \lambda)f(x_2).$$

A function is said to be strictly convex if the equality is strict for $x_1 \neq x_2$.

A concave function is simply the reversed inequality, such that $f(\lambda x_1 + (1-\lambda)x_2) \geq \lambda f(x_1) + (1 - \lambda)f(x_2)$.

Definition 2 (Preference relations)**.** *Let X be a set of outcomes, then for all $x, y \in X$:*

a. If $x \succ y$ then x is strictly preferred to y.

b. If $x \succeq y$ then x is weakly preferred to y.

c. If $x \sim y$ then x is indifferent to y.

A *utility* function measures an investor's relative risk preference for different levels of total wealth, and is a mapping that represents a weakly preferred preference relation. Typically, a utility function is a twice-differentiable function of wealth $U(w)$ defined for $w > 0$ which has the properties of non-satiation ($U'(w) > 0$) and risk aversion ($U''(w) < 0$). Utility functions also allow for positive affine transformations, in order to compare investments. There are three different risk categorisations based on the curve of the utility function; risk averse, risk neutral or risk seeking. These categories can be explained with the second derivative, or with Jensen's Inequality.

Theorem 1 (Jensen's Inequality)**.** *Let $f(x)$ be a convex function and X a random outcome, then*

$$\mathbb{E}\left[f(X)\right] \geq f(\mathbb{E}[X]).$$

Definition 3 (Risk preferences)**.** *Let $w \in \mathbb{R}^+$ and X be a positive random outcome.*

a. An investor is said to be risk averse if $U''(w) < 0$ or equivalently $\mathbb{E}[U(X)] < U(\mathbb{E}[X])$.

b. An investor is said to be risk neutral if $U''(w) = 0$ or equivalently $\mathbb{E}[U(X)] = U(\mathbb{E}[X])$.

c. An investor is said to be risk seeking if $U''(w) > 0$ or equivalently $\mathbb{E}[U(X)] > U(\mathbb{E}[X])$.

The thesis will focus on the risk averse preference. The intuition of a concave utility function is that for someone who only has one dollar to start with, obtaining one more dollar is quite important. For someone who already has a million dollars, obtaining one more dollar is nearly meaningless. The non-satiation property states that utility increases with wealth, i.e that more wealth is always preferred to less wealth, and that the investor is never satiated.

The *principle of expected utility maximisation* states that a rational investor, when faced with a choice among a set of competing feasible investment alternatives, acts to select an investment which maximises his expected utility of wealth. The four axioms that define this rational behaviour can now be stated. Let A, B and C be mutually exclusive random outcomes with known probabilities.

Axiom 1 (Completeness). *Exactly one of the following holds:* $A \prec B$, $A \succ B$ *or* $A \sim B$.

Axiom 2 (Transitivity). *Preferences are consistent across any three options: If* $A \succeq B$ *and* $B \succeq C$ *then* $A \succeq C$.

Axiom 3 (Independence). *If* $A \succeq B$ *and let* $p \in (0, 1]$, *then* $pA + (1-p)C \succeq pB + (1-p)C$.

The axiom states that outcomes mixed with a third one maintain the same preference order as when the two are presented independently of the third one.

Axiom 4 (Continuity). *If* $A \preceq B \preceq C$, *then there exists a probability* $p \in [0, 1]$ *such that* $pA + (1-p)C \sim B$.

The axiom assumes that there is a "tipping point" between being better or worse off than a given middle outcome.

If all these axioms are satisfied, then the individual is said to be rational, and the preferences can be represented by a utility function.

Theorem 2 (Utility function). *If preferences satisfy Axioms 1-4, then there exists a utility function* $U : X \to \mathbb{R}$ *such that for two random outcomes* X, Y *the relation* $X \succeq Y$ *holds if and only if* $\mathbb{E}[U(X)] \geq \mathbb{E}[U(Y)]$.

Such a function is referred to as a *von Neumann-Morgenstern utility.* For proof, please refer to the original work by Neumann and Morgenstern (1944).

The degree of risk aversion is determined by the curvature of the utility function. *Absolute risk aversion* is a measure of the investor's reaction to uncertainty in absolute (dollar) terms

$$ARA(w) = -\frac{U''(w)}{U'(w)}, \tag{2.1}$$

while the *Relative risk aversion* is a measure proportional to changes of wealth

$$RRA(w) = -\frac{wU''(w)}{U'(w)}. \tag{2.2}$$

A constant RRA implies a decreasing ARA, but the opposite is not always true. If an investor is characterised by decreasing (increasing) absolute risk aversion, it implies that the utility function is positively (negatively) skewed, and he/she will commit more (less) dollars to risky investments as wealth increase. A decreasing (increasing) relative risk aversion means he/she will commit a larger (smaller) fraction of wealth as wealth increases.

Now we can introduce some common utility function classes. Hyperbolic Absolute Risk Aversion is a family of utility functions that satisfies the four axioms, where the utility functions are of the following form.

Definition 4 (HARA Utility). *A utility function that exhibits hyperbolic absolute risk aversion is defined as*

$$U(w) = \frac{1-\gamma}{\gamma}\left(\frac{aw}{1-\gamma} + b\right)^{\gamma},$$

with $a \in \mathbb{R}_{>0}$, $\gamma \neq 1$, $\frac{aW}{1-\gamma} + b > 0$ *or* $b = 1$ *if* $\gamma = -\infty$.

The utility function exhibits hyperbolic absolute risk aversion if and only if the risk tolerance $T(w)$ is linear in wealth, where

$$T(w) = \frac{1}{ARA(w)} = \frac{w}{1-\gamma} + \frac{b}{a}. \tag{2.3}$$

The HARA utility function is commonly used due to the mathematical tractability, and special cases include quadratic utility functions, exponential utility functions, and iso-elastic utility functions.

Another utility function of interest is the Constant Relative Risk Aversion, where risk is proportional to wealth (dollar risk increases with wealth).

Definition 5 (CRRA Utility). *A utility function that exhibits constant relative risk aversion is defined as*

$$U(w) = b + a\frac{w^{\gamma}}{\gamma},$$

with $a, b \in \mathbb{R}$ *and* $\gamma \neq 0$. *If* $\gamma = 0$ *then* $U(w) = log(w)$ *is used instead.*

The HARA family nests CRRA when $\gamma < 1$ and $b = 0$.

Prospect theory, which is closely related to utility theory, explains why people make non-optimising decisions when presented with uncertain outcomes. The main difference is that people treat perceived losses and wins differently.

Definition 6 (Prospect function)**.** *A prospect function with preference homogenity has the power function form*

$$U(x) = \begin{cases} x^a, & \text{if } x \geq 0, \\ -c(-x)^b, & \text{otherwise,} \end{cases}$$

where $a > 0$, $b > 0$ *and* $c > 0$. *Loss aversion is implied if* $c > 1$ *or* $a = b$.

For an exposition of prospect theory, please refer to Kahneman and Tversky (1979).

The thesis will focus on an affine transformation of the HARA utility in the form

$$U(w) = \frac{(w-b)^\gamma}{\gamma}, \qquad \gamma < -1, \quad b \geq 0, \quad w > b.$$

The function has the properties $U' > 0$, $U'' < 0$ where both the absolute and relative risk aversion is decreasing in w (assuming $b > 0$), implying that both the dollar amount and fraction of wealth risked increases with wealth. If $b = 0$ it equals the CRRA model, and the relative risk aversion is therefore constant. If b is positive, the investor is very risk averse close to this subsistence level b which decreases as wealth increases.

2.2 Dynamic programming

The following defines the framework of a non-stationary finite horizon Markov Decision Model, which is the type of model used to define and solve the stochastic control problems in the thesis. The purpose is to optimise sequential decision-making under uncertainty. Such optimisation problems can be solved with backward induction algorithms, and this section describes the framework and under what assumptions optimal policies exist. The theory presented in this section is based on Bäuerle and Rieder (2011).

First, define the following spaces.

Definition 7 (State space). *A space E endowed with σ-algebra $\mathcal{E}$, is called a state space. The states are denoted by $x \in E$.*

Definition 8 (Action space). *A space A endowed with σ-algebra $\mathcal{A}$, is called an action space. The actions are denoted by $a \in A$.*

Definition 9 (Disturbance space). *A space $\mathcal{Z}$ endowed with σ-algebra $\mathfrak{Z}$, is called a disturbance space. A random disturbance is denoted by $Z \in \mathcal{Z}$, and the realisation by $z \in \mathcal{Z}$.*

An agent exists at time t in a given state x_t. The state follows a stochastic process $(X_t)_{t=0}^{T}$, which is controlled by an action a_t and affected by random disturbance Z_t. At any time $t = 0, ..., T-1$ the agent can take an *admissible action* $a \in \mathcal{A}$ to change the law of transition from t to $t+1$.

Definition 10 (Admissible action). *A measurable subset $D_t \subset E \times A$ denotes the set of possible state-action combinations at time t, where D_t contains the measurable mapping $f_t : E \to A$. The set $D_t(x) = \{a \in A | (x, a) \in D_t\}$ is called the set of admissible actions in state x at time t.*

The action results in an immediate reward $r_t(x_t, a_t)$ which depends on the current state x_t and action a_t. At the terminal time T the agent cannot take action, but receives terminal reward $g_T(x_T)$.

Definition 11 (Reward function). *A measurable function $r_t : D_t \to \mathbb{R}$ is called a reward function, where $r_t(x, a)$ gives the one-period reward of a system being in state x and action a is taken at time t.*

Definition 12 (Terminal reward function). *A measurable function $g_T : E \to \mathbb{R}$ is called a terminal reward function, where $g_T(x)$ gives the terminal reward of a system being in state x at time T.*

The evolution of the state variables over time is described by the *transition function* $T_t(x_t, a_t, z_{t+1})$, where the probability to reach a certain state x' at $t+1$ is given by the *stochastic kernel* $Q_t(dx' | (x_t, a_t))$.

Definition 13 (Transition function). *A measurable function $T_t : D_t \times \mathcal{Z} \to E$ is called a transition function, where $T_t(x, a, z)$ gives the transition of the system from state x at time t before action, if action a is taken and disturbance z occurs at time $t + 1$.*

Definition 14 (Stochastic transition kernel). *The mapping $B \mapsto Q_t(B|(x, a))$ is the probability to reach a state in $B \subset E$ at time $t + 1$ if the action $a \in A$ is applied to the state $x \in E$ at time t.*

The transition law for each period is therefore determined by T_t and Q_t. The Markov Decision Model can now be established.

Definition 15 (Markov Decision Model). *A (non-stationary) Markov Decision Model with finite planning horizon $T \in \mathbb{N}$ consists of a set of data $(E, A, D_t, \mathcal{Z}, T_t, Q_t, r_t, g_T)$ for $t = 0, 1, ..., T - 1$.*

The goal is to maximise the expected total reward accumulated from applying decision rules to the system during the time horizon $t \in \{0, ..., T\}$, by finding an optimal policy for the full period.

Definition 16 (Decision rules and policy). *A measurable mapping $\pi_t : E \to A$ with the property $\pi_t(x) \in D_t(x) \, \forall \, x \in E$ is called a decision rule at time t. The set of all decision rules at time t is denoted F_t. A sequence of decision rules $\pi = (\pi_0, \pi_1, ..., \pi_{T-1})$ with $\pi \in F_t$ is called a policy.*

Definition 17 (Expected total reward). *The expected total reward for state $x \in E$ at time t for the period $t, t + 1, ..., T$ if policy π is used is given by*

$$V_{t,\pi}(x) := \mathbb{E}_t^{\pi} \left[\sum_{k=t}^{T-1} r_k(X_k, \pi_k(X_k)) + g_T(X_T) \,\middle|\, X_t = x \right].$$

The maximum expected total reward is then given by $V_t(x) := \sup_\pi V_{t,\pi}(x)$, and often referred to simply as the *value function*, where the policy is the *optimal policy* if $V_{t,\pi}(x) = V_t(X)$ for all $x \in E$.

To ensure that all expectations are well-defined, the integrability assumption is introduced.

Definition 18. *[Integrability Assumption] The integrability assumption is satisfied if for $t = 0, 1, ..., T$ and $x \in E$,*

$$\sup_{\pi} V_{t,\pi}(x) < \infty.$$

A sufficient condition for the integrability assumption is the existence of an *upper bounding function.* For proof, and further arguments for sufficient conditions, please refer to Bäuerle and Rieder (2011) Section 2.4.

Definition 19 (Upper bounding function)**.** *A measurable function $b : E \to \mathbb{R}_+$ is called an upper bounding function with constant $C \in \mathbb{R}_+$ if for $x \in E$, $a \in A$, $t = 0, ..., T-1$ the following holds:*

$$|r_t(x,a)| \leq Cb(x), \quad |r_T(x)| \leq Cb(x), \quad \int_E b(x')Q_t(dx'|x,a)) \leq Cb(x)).$$

In order to find $V_t(x)$ we must also find $V_{t+1}(x)$. The problem therefore has a recursive solution. The calculation of the optimal policy is addressed in the following settings. Let $\mathbb{M}(E)$ denote the set where a function $v : E \to [-\infty, \infty)$ is measurable for the state space E. We have $V_{t,\pi} \in \mathbb{M}(E) \, \forall \pi, t$ due to the assumptions, as it is not generally true that V_t is measurable. For $t = 0, .., T-1$ and $v \in \mathbb{M}(E)$, define the operator

$$(L_t v)(x,a) = r_t(x,a) + \mathbb{E}\left[v(T_t(x_t, a_t, Z_{t+1}))\right], \quad (x,a) \in D_t.$$

With the operator L_t, define the reward operator (or so-called *Bellman operator*).

Definition 20 (Reward operator)**.** *For $t = 0, ..., T-1$, $v \in \mathbb{M}(E)$ and $\pi \in F_t$ define*

$$\mathcal{T}_{t,\pi} v(x) = (L_t v)(x, \pi(x)), \quad x \in E.$$

The *maximal reward operator* is then given by $\mathcal{T}_t v(x) = \sup_\pi \mathcal{T}_{t,\pi} v(x)$, and the associated policy π is a *maximiser* of v at time t if $\mathcal{T}_{t,\pi} v = \mathcal{T}_t v$. The Bellman operators are monotone, and will be used to compute the value function recursively with *Reward Iteration.*

Theorem 3 (Reward Iteration). *Let* $\pi = (\pi_0, ..., \pi_{T-1})$ *be a policy. For* $t = 0, 1, ..., T-1$ *it holds that*

a) $V_{T,\pi} = g_T$ *and* $V_{t,\pi} = \mathcal{T}_{t,\pi_t} V_{t+1,\pi}$.

b) $V_{t,\pi} = \mathcal{T}_{t,\pi_t} ... \mathcal{T}_{T-1,\pi_{T-1}} g_T$.

The sequence of the reward operators shall be interpreted as $\mathcal{T}_{t,\pi_t}(\mathcal{T}_{t+1,\pi_{t+1}}(...))$ where the inner operator is applied first. For the proof of the theorem, please refer to Bäuerle and Rieder (2011) p.21. A problem written in the form of point (a) in Theorem 3 is said to be a *Bellman equation.* Under appropriate assumptions, there exists a recursive solution $(v_t)_{t=0}^T$ to the Reward Iteration, which yields the value functions and determines an optimal policy. This is stated in the verification theorem, which is sufficient when state and action spaces are finite.

Theorem 4 (Verification Theorem). *Let* $(v_t) \subset \mathbb{M}(E)$ *be a solution of the Bellman equation. Then it holds:*

a) $v_t \geq V_t \quad \forall t$.

b) If π_t^* *is a maximiser of* v_{t+1} *for* $t = 0, 1, ..., T-1$ *then* $v_t = V_t$ *and the policy* $\pi^* = (\pi_t^*, ..., \pi_{T-1}^*)$ *is optimal for the Markov Decision Problem.*

For proof, please refer to Bäuerle and Rieder (2011) p.22.

The *Structure Theorem* verifies the necessary structure of the Markow Decision Problem in order for it to be solved with backward recursion.

Theorem 5. *Let the following structure assumptions be satisfied. There exist sets* $\mathbb{M}_t \subset \mathbb{M}(E)$ *and* $\Delta_t \subset F_t$ *such that for all t=0,1,...,T-1:*

(i) $g_T \in \mathbb{M}_T$.

(ii) If $v \in \mathbb{M}_{t+1}$ *then* $\mathcal{T}_t v$ *is well-defined and* $\mathcal{T}_t v \in \mathbb{M}_t$.

(iii) For all $v \in \mathbb{M}_{t+1}$ *there exists a maximiser* π_t *of* v *with* $\pi_t \in \Delta_t$.

Then it holds:

a) $V_t \in \mathbb{M}_t$ *and the sequence* (V_t) *satisfies the Bellman equation for* $t = 0, ..., T-1$

$$V_T(x) = g_T(x),$$

$$V_t(x) = \sup_{a \in D_t(x)} \{r_t(x,a) + \mathbb{E}[V_{t+1}(T_t(x_t, a_t, Z_{t+1}))]\}, \quad x \in E.$$

b) $V_t = \mathcal{T}_t \mathcal{T}_{t+1} ... \mathcal{T}_{T-1} g_T$.

c) For n=0,1,...,T-1 there exist maximisers π_t *of* V_{t+1} *with* $\pi_t \in F_t$*, and every sequence of maximisers* π^* *of* V_{t+1} *defines an optimal policy* $(\pi_t^*, ..., \pi_{T-1}^*)$ *for the Markov Decision Problem.*

For proof, please refer to Bäuerle and Rieder (2011) p.23. The theorem states that the maximisers yield an optimal strategy.

If the structure assumptions are satisfied, the problem can now be solved with a *Backward Induction Algorithm*.

Algorithm 1 Backward Induction Algorithm

for $t \leftarrow T, 0$ **do**
  if $t = T$ **then**
    for $x \in E$ **do**
      $V_T(x) := g_T(x)$
    end for
  else if $0 \leq t < T$ **then**
    for $x \in E$ **do**
      [Compute the maximiser]
      $\pi_t^*(x) = \arg\max_{a \in D_t(x)} \{r_t(x,a) + \mathbb{E}[V_{t+1}(T_t(x,a,Z_{t+1}))]\}$
      [Update the value function with optimal decision]
      $V_t(x) = r_t(x, \pi_t^*(x)) + \mathbb{E}[V_{t+1}(T_t(x, \pi_t^*(x), Z_{t+1}))]$
    end for
  end if
end for

Chapter 3

Expected utility model for retirement

3.1 Introduction

Australians face a number of challenges in retirement. After accumulating Superannuation assets throughout their working lives, they now have access to all these funds, thus the responsibility for managing said funds changes from investment decisions only to include drawdown and consumption decisions. Not only do they face the risk of making poor consumption decisions, but they are also exposed to mortality and longevity risk, financial risks such as investment returns and inflation, as well as policy risk from changes to the Age Pension policies. All of these factors make it very difficult to plan ahead in retirement, and both retirees and advisors have limited knowledge regarding how to best consume and manage these assets.

The consumption choice itself can have severe consequences. If assets are consumed too quickly in retirement and the funds are depleted, then the retiree must rely on Age Pension payments for the remainder of his or her life. If consumption is instead kept to a minimum, the retirees will miss out on utility from consumption and might miss out on 'free' money from Age Pension payments. This results in confusion for many retirees (Agnew et al., 2013). Once retired, Australian retirees tend to keep the same proportion of risky assets as before retirement, even though the exposure decreases with age (Spicer et al., 2016). With no labour income, these assets are subject to sequencing risk[1], which is difficult to avoid in the decumulation phase, especially early in retirement (Kingston and Fisher, 2014). Conventional recommendations for asset management are no longer applicable in this case, as they are highly dependent on wealth level, age, and Age Pension policies.

Empirical data indicate that 75% of retirees are homeowners, whose homes account for 80% of the total wealth in the middle-wealth quartiles and nearly all of the wealth in the lower quartile[2]. Consumption tends to decrease with age and converges towards a consumption floor. It is argued that this decline in consumption is due to the declining health of retirees, which reduces their capacity for activity (Clare, 2014; Yogo, 2016), or due to retirees having fewer resources reserved for consumption owing to longevity risk (Milevsky and Huang, 2011). Higgins and Roberts (2011) and Asher et al. (2017) found that the rate of decline with age is also dependent on wealth, and expenditures tend to converge towards a constant level as the retiree ages. A model that properly captures these characteristics is required to estimate the wealth needed in retirement and to forecast Age Pension budgets and policy changes from a government perspective.

There have been attempts to model the decumulation phase with respect to the Australian Age Pension, such as Ding (2014), Hulley et al. (2013) and Iskhakov et al. (2015). These models are rather limited, however, as they either have a certain focus, or require an analytical solution. For an analytical solution to exist, it constrains either the number of stochastic variables allowed or the structure and assumptions of the model. Ding (2014) finds a semi-analytical solution to the problem, but imposes a sequential assumption where the retiree goes through the stages of no Age Pension, partial due to the asset-test, partial due to the income-test and finally full Age Pension. This order is often violated with stochastic returns. In addition, Ding (2014) does not include mortality risk other than as a weight for decreasing consumption with age. While Hulley et al. (2013) solve the model numerically, the model is based on CRRA utility which does not capture wealth dependent preferences among retirees and does not allow for a consumption floor

[1]Sequencing risk refers to the unfortunate timing of cash flow in the investment portfolio. Drawdown after negative investment returns will have a greater impact on the long-term growth of the portfolio than drawdown after positive investment returns.

[2]Estimated from data (Australian Department of Treasury, 2010) but consistent with Olsberg and Winters (2005).

to be defined. Iskhakov et al. (2015) apply HARA utility which does not have the same shortcomings as in Hulley et al. (2013), but focus solely on optimal annuity allocation. In addition to these, there are a few important extensions of traditional utility maximisation models from an Australian retirement perspective: the treatment of housing in utility models (Yogo, 2016; Cho and Sane, 2013), the definition of the bequest function (Lockwood, 2014; De Nardi, 2004; Ameriks et al., 2011) and how health affects expenditure endogenously in life cycle modelling (Yogo, 2016).

The contribution of this chapter is a sequential family status model in the retirement stage that considers stochastic wealth, stochastic family status, and mortality risk, as well as a "health" status proxy that allows the model to fit empirical data better. The focus is on the decumulation phase of the life cycle problem from the point in time when an individual retires. The model is set up as a stochastic dynamic programming problem and solved numerically via backward recursion.

The remainder of this chapter is organised as follows. Section 3.2 defines the model that includes the Age Pension means-test and presents the corresponding optimal stochastic control problem. Section 3.3 explains the numerical implementation to solve the model, and Section 3.4 presents and discusses the results. Finally, Section 3.5 concludes the chapter.

3.2 Model specification

We assume that the agent's goal is to maximise the expected value of utility associated with consumption, housing, and bequest. Utility is measured by time-separable additive functions based on commonly used HARA utility functions.

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$ be a filtered complete probability space, and let $\mathcal{F}_t$ represent the information available up to time t. We assume that all the processes introduced below are well defined and adapted to $\{\mathcal{F}_t\}_{t\geq 0}$. Denote the value of liquid financial assets as random variable W_t and family status as random variable G_t at the agent's anniversary dates $t = t_0, t_0 + 1, ..., T$, where t_0 is the retirement age and T is the maximum age of the agent beyond which survival is deemed impossible. Realisations of W_t and G_t are denoted as w_t and g_t, respectively. The utility received at times t is subject to the agent decision (control) variables α_t (proportion drawdown of liquid assets) and δ_t (proportion of liquid assets allocated to risky assets) as well as the decision variable ϱ (wealth allocated to total housing at time $t = t_0$). Given a current state $x_t = (w_t, g_t)$, we can define a decision rule $\pi_t(x_t) = (\alpha_t(x_t), \delta_t(x_t))$, which is the action at time t, and x_t is the value of the state variables before the action. Then, a sequence (policy) of decision rules is given by $\pi = (\pi_{t_0}, \pi_{t_0+1}, ..., \pi_{T-1})$ for $t = t_0, t_0 + 1, ..., T - 1$.

The optimisation problem is set to begin at the time of retirement $t = t_0$ when all available wealth is placed in an Allocated Pension account, and the agent decides how much of the total wealth will be allocated to a family home[3]. Taxes after t_0 are not considered because earnings in an Allocated Pension account are tax-free. At the start of each year, t, the agent makes a decision as to how much to withdraw from the account, and receives a means-tested Age Pension P_t. The remaining wealth is placed in a stochastic risky portfolio S_t and risk-free cash account until $t+1$, where the agent has to decide what proportion δ_t has to be allocated to risky assets. Then, the wealth process and consumption are given by

$$W_{t+1} = [W_t - \alpha_t W_t] \left[\delta_t e^{Z_{t+1}} + (1 - \delta_t) e^{r_t}\right], \quad W_{t_0} = \mathsf{W} - H, \tag{3.1}$$

$$H = \varrho \mathsf{W}, \tag{3.2}$$

$$C_t = \alpha_t W_t + P_t, \tag{3.3}$$

where $W_t \geq 0$ denotes the liquid assets at time t before withdrawal. The initial liquid assets W_{t_0} constitute the remaining wealth after housing allocation at retirement, where W is the total wealth[4] at time $t = t_0$. The housing allocation is constrained by $\varrho \in \{0, [\frac{H_L}{\mathsf{W}}, 1]\}$, where $H_L \geq 0$ is a lower bound of housing; hence, housing wealth can never be negative. The allocation, if non-zero, must therefore be larger than this lower bound. It is assumed that the risky asset S_t follows a geometric Brownian motion such that the real log returns of the risky asset $Z_{t+1} = \ln(S_{t+1}/S_t)$,

[3]Note that H represents the market value of the house at t_0 and not additional wealth invested into housing. As the retiree most likely is a homeowner already, the difference between H and the current house value represents the suggested change in housing. A more realistic assumption would be to allow the retiree to scale down housing later in retirement, but owing to limitations in the dataset, such a model cannot be calibrated.

[4]Total wealth includes current house value if the retiree is a homeowner; thus, allocation of part of the initial wealth to housing H corresponds to up-/downscaling the house depending on the current house value.

$t = 0, 1, \ldots$ are independent and identically distributed from a normal distribution $\mathcal{N}(\mu - \widetilde{r}, \sigma^2)$ with mean $\mu - \widetilde{r}$ and variance σ^2, where μ is the mean risky return and $\widetilde{r}$ is the inflation rate. The real risk-free rate r_t (adjusted for inflation) is time-dependent but deterministic. The constraints for consumption, where C_t is the consumption and P_t is the Age Pension payments, indicate that consumption equals the sum of received Age Pension and drawdown of wealth for the current period. By defining the model in real terms (adjusted for inflation), we can avoid dealing with inflation in consumption, the Age Pension, and the consumption floor to capture the characteristics of Australian retirees more effectively.

The model operates at a household level, which distinguishes between couple and single retiree households because the Age Pension treats couples as a single entity. The agents face the risk of family status transitions due to death in each period, with the possible states defined as

$$G_t \in \mathcal{G} = \{\Delta, 0, 1, 2\}, \tag{3.4}$$

where Δ corresponds to the agent already being dead at time t, 0 corresponds to the death of the agent during $(t-1, t]$, 1 corresponds to the agent being alive at time t in a single household, and 2 corresponds to the agent being alive in a couple household, subject to survival probabilities. Therefore, an agent can start at time $t = t_0$ as either a couple or a single household. In the case of a couple household, there is a risk in each time period that one spouse will pass away, in which case it is treated as a single household model for the remaining years[5]. Further, Z_t and G_t are assumed to be independent; hence, a large investment loss does not affect the death probabilities (although one can argue that it might affect, e.g., the quality of life or the ability to pay hospital bills, which in turn would affect death probabilities).

[5]While it is possible for a retiree to form new relationships in retirement, the proportion of total marriages for those aged 65-69 was 0.1% for females and 0.2% for males in 2014 (Australian Bureau of Statistics, 2015)—a proportion that only decreases with age. Even if partnerships without marriage are more likely in retirement, the overall probability is low and will not be included.

In each period, the agent receives utility based on the current family status G_t:

$$R_t(W_t, G_t, \alpha_t, H) = \begin{cases} U_C(C_t, G_t, t) + U_H(H, G_t), & \text{if } G_t = 1, 2, \\ U_B(W_t, H), & \text{if } G_t = 0, \\ 0, & \text{if } G_t = \Delta. \end{cases} \tag{3.5}$$

Thus, if the agent is alive, he/she receives a reward based on consumption U_C and housing U_H. If he/she died during the year, the reward comes from the bequest U_B, and if he/she is already dead, there is no reward. Note that the reward received when the agent is alive depends on whether the state is a couple or single household owing to differing utility parameters and Age Pension thresholds. The terminal reward function at $t = T$ is given by

$$\widetilde{R}(W_T, G_T, H) = \begin{cases} U_B(W_T, H), & \text{if } G_T \geq 0, \\ 0, & \text{if } G_T = \Delta. \end{cases} \tag{3.6}$$

The current model is based on the work of Ding (2014), but with the addition of sequential family status, mortality risk, and risky asset, and a different "health" proxy. The retiree wants to find the policy that maximises the expected utility with respect to their decision for consumption, investment, and housing. This is defined as a stochastic control problem:

$$\widetilde{V} := \max_{\varrho} \left[\sup_{\pi} \mathbb{E}^{\pi}_{t_0} \left[\beta_{t_0,T} \widetilde{R}(W_T, G_T, H) + \sum_{t=t_0}^{T-1} \beta_{t_0,t} R_t(W_t, G_t, \alpha_t, H) \right] \right], \tag{3.7}$$

which can be solved with dynamic programming by using backward induction of the Bellman equation. Further, $\mathbb{E}^{\pi}_{t_0}[\cdot]$ is the expectation conditional on information at time $t = t_0$ if policy π is used up to $t = T - 1$. The policy π includes the control variables for each time period, and $\beta_{t,t'}$ is the discount from t to t'.

The subjective discount rate $\beta_{t,t'}$ is a proxy for personal impatience between time t and t', and it is set in relation to the real interest rate such that

$$\beta_{t,t'} = e^{-\sum_{i=t}^{t'} r_i}. \tag{3.8}$$

This assumption suggests that optimal consumption rates would be constant over time for the HARA utility in the absence of mortality risk and risky investments[6].

3.2.1 Consumption preferences

We assume that the HARA utility comes from consumption exceeding the consumption floor, weighted with a time-dependent "health" status proxy. The "health" proxy does not intend to model the retiree's health, as the dataset does not contain such information, but is rather a means of explaining decreasing consumption. The utility function for consumption is defined as

$$U_C(C_t, G_t, t) = \frac{1}{\psi^{t-t_0}\gamma_d}\left(\frac{C_t - \overline{c}_d}{\zeta_d}\right)^{\gamma_d}, \quad d = \begin{cases} \text{C}, & \text{if } G_t = 2 \quad \text{(couple)}, \\ \text{S}, & \text{if } G_t = 1 \quad \text{(single)}, \end{cases} \tag{3.9}$$

where $\gamma_d \in (-\infty, 0)$ denotes the risk aversion and $\overline{c}_d$ is the consumption floor parameter. The scaling factor ζ_d normalises the utility that a couple receives in relation to a single household[7]. The utility parameters γ_d, $\overline{c}_d$, and ζ_d are subject to family state G_t and will have different values for couple and single households. Note that the consumption for any period is based on the pension P_t received in the same period and the drawdown α_t from the liquid assets W_t, as given by equation (3.3). The proportion of wealth drawn down can be positive or negative, with a negative value indicating that part of the pension received is saved for future consumption. Consumption tends to converge towards a consumption floor as the retiree ages despite their wealth status: hence, we define a "health" proxy to control the slope of the decline. Let $\psi \in [1, \infty)$ be the utility parameter for the slope, where the difference between the current time t and the time of retirement t_0 determines the power of the parameter. This prevents the initial time $t = t_0$ from being affected by a "health" proxy, and as the retiree ages, the slope of the proxy decreases, allowing the consumption to converge over age and wealth groups. This decreasing convex "health" proxy has a better fit to empirical data compared with survival probabilities as used by Ding (2014), which are decreasing but concave. An alternative would be to utilise

[6] This can easily be shown with simple calculus.

[7] If single and couple households had the same risk aversion and no scaling factor was used, the solution would suggest similar consumptions for both. This effect comes from the consumption smoothing properties of life cycle models and thus needs to be adjusted, as a couple household's utility is for two people. Otherwise, it would cause problems in the calibration stage.

a non-rational model, rather than a rational model with the extra utility parameter, such as in Marín-Solano and Navas (2010).

3.2.2 Housing preferences

Housing differs from other assets in that it provides a flow of services in terms of the preference (utility) of owning a house compared to renting, in addition to its residual value. The reasons for including housing in the model are two-fold. First, housing allocation can be used as a means of Age Pension planning (leaving less wealth so the retiree is entitled to more Age Pension). Second, allocation to housing implies the sacrifice of future consumption owing to less liquid wealth being available, thereby affecting consumption preferences. In addition, the model allows for house up-/downscaling at retirement time t_0. We apply the assumption that the utility is linked to the house value as in Ding (2014) and Cho and Sane (2013). The utility from owning a home is defined as

$$U_H(H, G_t) = \begin{cases} \frac{1}{\gamma_{\mathrm{H}}} \left(\frac{\lambda_d H}{\zeta_d} \right)^{\gamma_{\mathrm{H}}}, & \text{if } H \geq 0, \\ 0, & \text{otherwise,} \end{cases} \tag{3.10}$$

where γ_{H} is the risk aversion parameter for housing (different from risk aversion for consumption and bequest), ζ_d is the same scaling factor as in equation (3.10), H is the market value of the family home at the time of purchase t_0, and $\lambda_d \in (0, 1]$ is the housing preference defined as a proportion of the market value.

Note that the house value H is not indexed with time. The retiree decides how much of his/her total wealth to allocate to housing at the time of retirement t_0, which remains constant (in real terms) afterwards. If the retiree already is a homeowner, H represents the target housing allocation after up- or downscaling. We do not consider the house to be a liquid asset but a proxy for the utility received by a homeowner. Our assumptions reflect the housing behaviour of Australian retirees. Most Australian households do not convert housing assets into liquid assets in order to cover expenses in retirement, with the exception of certain events (such as the death of a spouse, divorce, or moving to an aged care facility) (Olsberg and Winters, 2005; Asher et al., 2017). The dataset does not include such information: hence, a simplistic assumption is made, as we cannot calibrate a model that allows for the option of changing housing after retirement. Housing also tends to be constant over age groups (Ding, 2014): hence, the option of scaling down *during* retirement is not required to model Australian retirement. Wealthier retirees prefer to invest more in the family home as their wealth increases, but with a decreasing marginal utility because the percentage allocation decreases, which is consistent with the utility model used.

3.2.3 Bequest preferences

We adopt the bequest function in Lockwood (2014), which is a re-parameterised version of De Nardi (2004), as the parameters are slightly more intuitive. The utility is received from both remaining liquid wealth and the residual value of the home. The home is not considered to represent a growth yielding asset, however, as housing is both a necessity and an asset. In addition to the arguments in Section 3.2.2, the intentional bequest component is difficult to separate, and housing is not considered for bequest purposes at the time of purchase (Olsberg and Winters, 2005). In addition, since the house market value is used as a proxy for utility derived from owning a house in terms of the service it provides, an increasing house value cannot suggest increased utility. Since the model does not allow for downscaling during retirement (hence housing assets cannot be accessed for consumption), modelling the home value as an increasing asset will only affect the preference between consumption and bequest, but not affect the level of consumption in relation to wealth. Then, the utility function is defined as

$$U_B(W_t, H) = \left(\frac{\theta}{1-\theta}\right)^{1-\gamma_S} \frac{\left(\frac{\theta}{1-\theta}a + W_t + H\right)^{\gamma_S}}{\gamma_S}, \tag{3.11}$$

where W_t denotes the liquid assets available for bequest, γ_S denotes the risk aversion parameters of bequest utility (which is considered to be the same as consumption risk aversion for singles, because a couple is expected to become a single household before bequeathing assets)[8], and the two parameters $\theta \in [0, 1)$ and $a \in \mathbb{R}^+$. The threshold for luxury bequest, a, is the threshold up to which the retiree leaves no

[8]In case couple households have a different risk aversion towards bequest, it will be absorbed by adjusting the ratio of single and couple risk aversion.

bequest[9]. If $a = 0$, then consumption and bequest are homothetic. The degree of altruism, θ, controls the preference of bequest over consumption. Low values of θ indicate that retirees prefer consumption to bequest, while high values increase the marginal utility of bequest. As $\theta \to 1$, the bequest motive approaches a linear function with a constant marginal utility of $\frac{dU_B}{dW_t} = a^{\gamma-1}$.

3.2.4 Age Pension function

The Age Pension received is subject to the asset and income means-test. This function is modelled with respect to the current liquid assets, where the account value is used for the asset-test. All liquid assets are converted into an Allocated Pension account at the time of retirement as per the model assumptions. This type of account has the advantage that earnings on assets are tax-free. The asset-test function is then defined as

$$P_{\text{A}} := P^d_{\max} - \left(W_t - L^{d,h}_{\text{A}}\right) \varpi^d_{\text{A}}, \tag{3.12}$$

where L^d_{A} is the threshold for the asset-test, ϖ^d_{A} is the taper rate for assets exceeding the thresholds, and superscript d is a categorical index indicating couple or single household status as defined in equation (3.10). The variables are subject to whether the household is a single or couple household, and the threshold for the asset-test is also subject to whether the retiree is a homeowner or not $h = \{0, 1\}$.

Since the model assumption states that no labour income is possible, all income for the income-test is either generated from withdrawals of liquid assets (accounts opened prior to the 1st of January 2015) or from deemed income (accounts opened after the 1st of January 2015). An Allocated Pension account, under the pre-2015 rules, allows for a yearly income-test deduction which is not available for new accounts. Instead, the new rules have introduced a 'work bonus' deduction for the income-test, but as the model assumes the retiree is no longer in the workforce this

[9] There is strong empirical evidence that wealthy retirees leave a larger proportion of their wealth as bequest compared with less wealthy retirees (Ameriks et al. (2011), Hurd and Smith (2003), Ding (2014)).

has been omitted in the function. The pre-2015 income-test function is defined as

$$P_{\mathrm{I}} := P_{\max}^{d} - \left(\alpha_t W_t - M(t) - L_{\mathrm{I}}^{d}\right) \varpi_{\mathrm{I}}^{d}, \tag{3.13}$$

where L_{I}^{d} and ϖ_{I}^{d} is the threshold and taper rate respectively with parameters for the income-test. The function $M(t)$ is an income-test deduction set when the wealth is converted into an Allocated Pension account, defined as

$$M(t) = \frac{W_{t_0}}{e_{t_0}}(1+\widetilde{r})^{t_0-t}, \tag{3.14}$$

where e_{t_0} is the expected lifetime at age t_0 and $\widetilde{r}$ is the inflation. As the model is defined in real terms, the future income-test deductions must discount inflation. The income-test for post-2015 accounts are based on deemed income rather than drawdowns, and can be written as

$$P_{\mathrm{I}} := P_{\max}^{d} - \left(P_{\mathrm{D}}(W_t) - L_{\mathrm{I}}^{d}\right) \varpi_{\mathrm{I}}^{d}, \tag{3.15}$$

$$P_{\mathrm{D}}(W_t) = \varsigma_{-} \min\left[W_t, \kappa^d\right] + \varsigma_{+} \max\left[0, W_t - \kappa^d\right]. \tag{3.16}$$

Here, $P_{\mathrm{D}}(W_t)$ calculates the deemed income, where κ_d is the deeming threshold, and ς_{-} and ς_{+} are the deeming rates that apply to assets below and above the deeming threshold respectively. The total Age Pension received can then be defined as

$$P_t := f(\alpha_t, W_t, t) = \max\left[0, \min\left[P_{\max}^{d}, \min\left[P_{\mathrm{A}}, P_{\mathrm{I}}\right]\right]\right], \tag{3.17}$$

where $P_{\max}^{d}$ is the full Age Pension. The Age Pension rules state that the entitlement age is 65 for both males and females, with the means-test thresholds and taper rates for July 2016 presented in Table 3.1.

3.2.5 Solution as a stochastic control problem

To solve the optimal stochastic problem of maximising the expected utility with respect to the decision policy, given by equation (3.7), the problem is defined as a dynamic programming problem (see Section 2.2).

The starting point is the basic model where the wealth W_t and family status G_t

Table 3.1: Age Pension rates published by Centrelink as at September 2016.

		Single	Couple
$P^d_{\max}$	Full Age Pension per annum	\$22,721	\$34,252
	Income-Test		
L^d_{I}	Threshold	\$4,264	\$7,592
ϖ^d_{I}	Rate of Reduction	\$0.5	\$0.5
	Asset-Test		
$L^{d,h=1}_{\mathrm{I}}$	Threshold: Homeowners	\$209,000	\$296,500
$L^{d,h=0}_{\mathrm{I}}$	Threshold: Non-homeowners	\$360,500	\$448,000
ϖ^d_{A}	Rate of Reduction	\$0.039	\$0.039
	Deeming Income		
κ^d	Deeming Threshold	\$49,200	\$81,600
ς_-	Deeming Rate below κ^d	1.75%	1.75%
ς_+	Deeming Rate above κ^d	3.25%	3.25%

are stochastic, and the terminal time T (time beyond which survival is deemed to be impossible) is fixed. The problem is defined as follows:

- Denote a state vector as $X_t = (W_t, G_t, H) \in \mathcal{W} \times \mathcal{G} \times \mathcal{H}$, where $W_t \in \mathcal{W} = \mathbb{R}^+$ denotes the current level of wealth and $G_t \in \mathcal{G} = \{\Delta, 0, 1, 2\}$ denotes whether the agent is dead, died in this period, is alive in a single household, or is alive in a couple household. The stages are sequential; hence, an agent that starts out as a couple becomes single when one spouse dies. $H \in \mathcal{H} = \mathbb{R}^+$ denotes wealth invested in housing at t_0, hence is not indexed by time.

- Denote an action space of $(\alpha_t, \delta_t, \varrho) \in \mathcal{A} = (-\infty, 1] \times [0, 1] \times \{0, [\frac{H_L}{\mathsf{W}}, 1]\}$ for $t = t_0$, and $(\alpha_t, \delta_t) \in \mathcal{A} = (-\infty, 1] \times [0, 1]$ for $t = t_1, ..., T-1$. Here, $\varrho \in \{0, [\frac{H_L}{\mathsf{W}}, 1]\}$ is the proportion of total wealth allocated to housing, $\alpha_t \in (-\infty, 1]$ denotes the proportion of wealth consumed and $\delta_t \in [0, 1]$ is the percentage of wealth allocated in the risky asset. The upper boundary of 1 indicates that the drawdown cannot be larger than our total wealth, nor can we invest more than 100% in risky assets; hence, borrowing is not allowed. However, negative values for drawdown are allowed as they represent savings from the Age Pension into wealth.

- Denote an admissible space of state-action combination as $D_t(x_t) = \{\pi_t(x_t) \in \mathcal{A} \mid \alpha_t \geq \frac{\bar{c}_d - P_t}{W_t}\}$, which includes the possible actions for the current state and indicates that withdrawals must be sufficiently large to cover the consumption floor.

- The transition function $T_t(W_t, \alpha_t, \delta_t, z_{t+1}) := W_{t+1} = W_t(1-\alpha_t) \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_t})$, where z_{t+1} is the realisation of the log return on the stochastic investment portfolio over $(t, t+1]$. We assume that the agent is small and cannot influence the asset price.

- Denote the stochastic transitional kernel as $Q_t(dx'|x, \pi_t(x))$, which represents the probability of reaching a state in $dx' = (dw_{t+1}, g_{t+1})$ at time $t+1$ if action $\pi_t(x)$ is applied in state x at time t. Since the transition function is based on the stochastic risky return Z_{t+1}, which is Markovian, the transition probability for wealth W_{t+1} is determined by the distribution of the risky return, where $Z_{t+1} \overset{i.i.d}{\sim} \mathcal{N}(\mu - \widetilde{r}, \sigma^2)$ with the probability density function denoted as $f_{\mathcal{N}}(z)$. Let $q(g_{t+1}, g_t)$ denote $\Pr[G_{t+1} = g_{t+1} \mid G_t = g_t]$. Since both state variables depend on exogenous and independent probabilities, we have

$$\begin{aligned}
&Q_t(dx'|x, \pi_t(x)) \\
&\quad = \Pr\left[W_{t+1} \in dw_{t+1}, G_{t+1} = g_{t+1} \mid X_t = x_t\right] \\
&\quad = \Pr[T_t(W_t, \alpha_t, \delta_t, Z_{t+1}) \in dw_{t+1}, G_{t+1} = g_{t+1} \mid W_t = w_t, G_t = g_t] \\
&\quad = \Pr[T_t(W_t, \alpha_t, \delta_t, Z_{t+1}) \in dw_{t+1} \mid W_t = w_t] \times q(g_{t+1}, g_t).
\end{aligned} \tag{3.18}$$

The probabilities for family status are defined as

$$\begin{aligned}
&q(2,2) = p_t^{\mathrm{C}}, \quad q(1,2) = 1 - p_t^{\mathrm{C}}, \\
&q(1,1) = p_t^{\mathrm{S}}, \quad q(0,1) = 1 - p_t^{\mathrm{S}}, \\
&q(\Delta, 0) = q(\Delta, \Delta) = 1,
\end{aligned} \tag{3.19}$$

where p_t^{C} is the probability of surviving for one more year as a couple or p_t^{S} as a single. All other transition probabilities for family status are 0.

- The reward function depends on the G_t state as defined in equation (3.5). If the agent is alive, he/she receives a reward based on consumption. If he/she died during the year, the reward comes from bequest, and if he/she is already dead, there is no reward. Note that the reward when the agent is alive depends on the Age Pension received and the consumption floor, which differs for couples and singles.

- The terminal reward function is defined in equation (3.6).
- The discount factor has been defined in equation (3.8), with $\beta_{t,t+1} \in (0, 1]$.

The optimal value function can now be stated as in equation (3.21) at starting time t_0. A solution for the stochastic control problem is given by a backward recursion Bellman equation

$$V_T(X_T) = \widetilde{R}_T(W_T, G_T, H), \tag{3.20}$$

$$V_t(X_t) = \sup_{\pi_t(x_t) \in D_t(x_t)} \left\{ R_t(W_t, G_t, \alpha_t, H) + \beta_{t,t+1} \, \mathbb{E}_t^{\pi} \left[V_{t+1}(X_{t+1}) \mid X_t \right] \right\}, \tag{3.21}$$

where $\mathbb{E}_t^{\pi}[\cdot]$ is calculated using the stochastic transition kernel $Q(\cdot)$ given in equation (3.18) as

$$\mathbb{E}_t^{\pi} \left[V_{t+1}(X_{t+1}) \mid X_t \right] = \sum_{g_{t+1} \in \mathcal{G}} \int_{-\infty}^{\infty} V_{t+1}(T_t(W_t, \alpha_t, \delta_t, z_{t+1}), g_{t+1}, H) f_{\mathcal{N}}(z_{t+1}) dz \times q(g_{t+1}, g_t). \tag{3.22}$$

The validity of the problem setup and existence of optimal policies is implied by the integrability assumption (Definition 18) and structure assumption (Theorem 5). The integrability assumption holds if the reward and terminal function are bounded from above, and if it satisfies $\mathbb{E}[Z_t] < \infty$ for all t where Z_t is the disturbance term. A power utility function with $\gamma < 0$ has an upper bound of 0, and a log-normal random variable has a finite expected value: hence, the integrability assumption is satisfied.

3.3 Numerical implementation

The Bellman equation for the value function V_t is solved recursively with backward induction by discretising the wealth and house asset state on a grid of log-equidistant grid points $W_0, ..., W_k$ and $H_0, ..., H_k$ for each year $t = t_0, ..., T$. The lower bound of the grid is set to \$1 as the utility for \$0 does not exist. The upper bound W_k is set to $W_k = \hat{W}_{\max} e^{(T-t_0)\mu + 5\sqrt{T-t_0}\sigma}$, where $\hat{W}_{\max}$ is the largest wealth to be modelled, in order to find a conservative upper bound based on the risky asset.

H_k is set to be twice the largest housing sample. This means that extrapolation has no material effect when integrating risky returns; hence, values close to the upper bound have no impact on the ranges $[W_0, \hat{W}_{\max}]$ and $[H_0, \hat{H}_{\max}]$ that are actually used in the solution. For each grid point in the wealth and house state, an optimal drawdown proportion α_t and risky asset allocation δ_t are found using a two-dimensional optimisation.

The value function V_t is interpolated between grid points based on the shape-preserving Piecewise Cubic Hermite Interpolation Polynomial (PCHIP) method, which preserves the monotonicity and concavity of the value function (Kahaner et al., 1989). The need to interpolate arises from the integration of the stochastic return. Since the value function is only available at predefined grid points, any values in between need to be interpolated. If traditional cubic splines were used, there would be a high probability of the function overshooting a point; hence, the solution could return a local maximum rather than a global maximum. Linear interpolation requires a much higher grid density at lower values owing to the steep derivative in these regions; hence, PCHIP is preferred. In general, given a function $f(x)$ and state x between two grid points $x_k \leq x \leq x_{k+1}$, the interpolant $P(x)$ of $f(x)$ for the kth interval is calculated as

$$\begin{aligned} P(x) =& \frac{3h_k(x-x_k)^2 - 2(x-x_k)^3}{h_k^3} f(x_{k+1}) + \frac{h_k^3 - 3h_k(x-x_k)^2 + 2(x-x_k)^3}{h_k^3} f(x_k) \\ &+ \frac{(x-x_k)^2(x-x_k-h_x)}{h_k^2} d_{k+1} + \frac{(x-x_k)(x-x_k-h_k)^2}{h_k^2} d_k, \end{aligned} \tag{3.23}$$

where $h_k = x_{k+1} - x_k$ and the slope d of the interpolant depends on the first divided difference $\Delta_k = (f(x_{k+1}) - f(x_k))/h_k$. If Δ_k and Δ_{k-1} are of opposite polarity, then $d_k = 0$; otherwise, d_k is given by

$$\frac{3h_k + 3_{k-1}}{d_k} = \frac{2h_k + h_{k-1}}{\Delta_{k-1}} + \frac{h_k + 2h_{k-1}}{\Delta_k}. \tag{3.24}$$

In addition, the conditions $P(x_k) = f(x_k)$, $P(x_{k+1}) = f(x_{k+1})$, $P'(x_k) = d_k$, and $P'(x_{k+1}) = d_{k+1}$ must hold.

The expectation, involved in the Bellman equation, with respect to the normally distributed stochastic return is calculated with Gauss-Hermite quadrature

$$\int_{-\infty}^{\infty} e^{-x^2} f(x)dx \approx \sum_{i=1}^{M} w(x_i)f(x_i), \tag{3.25}$$

where $w(x_i)$ is the weight and x_i is the node at which the value function is evaluated, which gives an exact result if $f(x)$ is any polynomial up to the order $2M-1$. For details on how to find the weights and nodes, we refer to the literature on numerical integration, e.g., Kahaner et al. (1989). The expectation is then calculated as

$$\begin{aligned} &\int_{-\infty}^{\infty} V_{t+1}\left(T_t(W_t, \alpha_t, \delta_t, z), G_{t+1}\right) f_{\mathcal{N}}(z)\, dz \\ &\quad = \int_{-\infty}^{\infty} \frac{e^{-x^2}}{\sqrt{\pi}} V_{t+1}\left(T_t\left(W_t, \alpha_t, \delta_t, \sqrt{2}\sigma x + \mu\right), G_{t+1}\right)\, dx \\ &\quad \approx \sum_{i=1}^{M} \frac{w(x_i)}{\sqrt{\pi}}\, V_{t+1}\left(T_t\left(W_t, \alpha_t, \delta_t, \sqrt{2}\sigma x_i + \mu\right), G_{t+1}\right), \end{aligned} \tag{3.26}$$

using $M = 5$ nodes. To speed up the calculations, a temporary vector is created each year, where the expectation of the value function is calculated for each grid point. By doing this prior to finding the optimal control, we avoid repeating the numerical integration during the optimisation as it is now sufficient with interpolation.

Finally, once the backward induction has reached $t = t_0$, the initial wealth W_0 can be determined by optimising the allocation between housing and liquid assets. The optimal path for a retiree can then be derived by following the optimal drawdown and risky allocation from the policy that corresponds to the wealth grid point for each time t, and repeated until the terminal condition at $t = T$.

Tests were performed with additional nodes for the Gauss-Hermite quadrature and larger bounds for W_k and H_k to verify the accuracy of the numerical solution. Solving the Bellman equation with 5, 10, or 25 nodes resulted in negligible differences; hence, five nodes were chosen to reduce the calculation time. In addition, a forward Monte Carlo simulation with random policies was generated to verify the optimality of the solution.

3.4 Model characteristic

The model is able to capture the characteristics of Australian retirees in terms of consumption, housing and wealth. However, no attempts are made to explain the reason for the empirical behaviour with the model. The model is solved with the arbitrary but realistic[10] parameters in Table 3.2, with Age Pension parameters taken from Table 3.1. The optimal decisions over time and wealth of single and couple households are almost identical, with the exception of actual dollar amounts, hence only singles are shown in this section.

Table 3.2: Parameters used for the solution.

γ_{S}	γ_{C}	γ_{H}	θ	a	$\bar{c}_{\mathrm{S}}$	$\bar{c}_{\mathrm{C}}$	ψ	λ	μ	σ	r	$\widetilde{r}$
-3	- 3	-3	0.95	20 000	10 000	15 000	1.1	0.03	0.08	0.1	0.01	0.03

Figure 3.1 shows that consumption increases with liquid wealth, but tapers off at higher wealth rather than a linear relationship. The consumption also decreases with age towards the consumption floor $\bar{c}_d$ regardless of wealth. The small irregularities seen in the otherwise smooth consumption surface appear when the Age Pension function is binding, hence consumption changes depending on whether no/partial/full Age Pension is received.

[10]These parameters are consistent with the calibration in the next chapter.

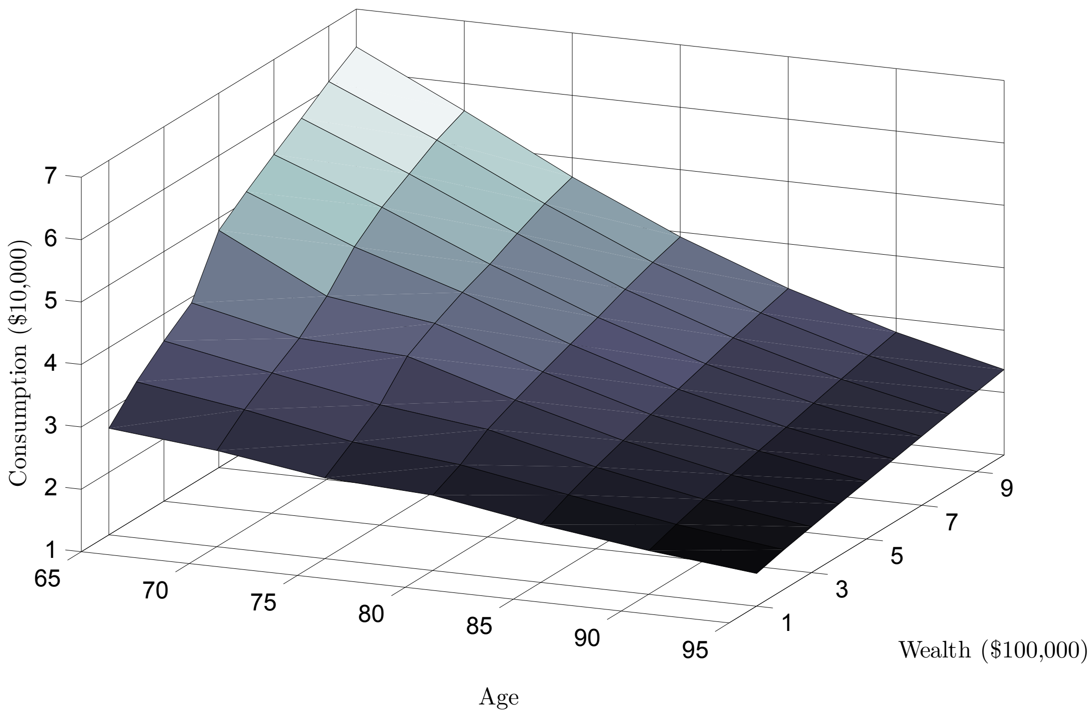

Figure 3.1: Optimal consumption given liquid wealth W_t and age, for a single non-homeowner household.

The proportion of total wealth that should be invested into housing (Figure 3.2) shows that for lower levels of total wealth, almost full allocation into housing is optimal. As total wealth increases, the proportion allocated to housing decreases, however, it will still remain high as the dollar amount allocated to housing increases monotonically with additional wealth. Very poor households that cannot cover the threshold for housing (e.g. the down payment) will not allocate any wealth to housing, which is shown by the kink in the bottom left of the curve.

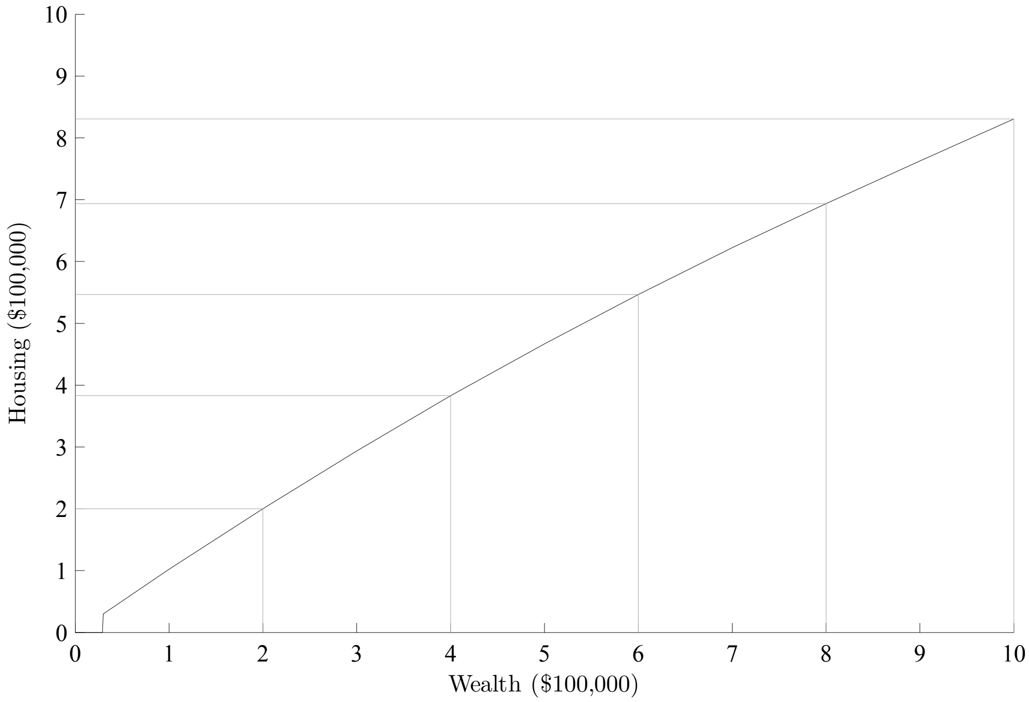

Figure 3.2: Optimal allocation to housing given total wealth W at time of retirement for a couple household.

Finally, retirees tend to preserve financial and residential wealth, consume conservatively, and pass on substantial bequest (Asher et al., 2017). While younger households decumulate assets, older ones accumulate. These characteristics can be seen in Figure 3.3, which shows the wealth paths throughout retirement for three different initial liquid wealth scenarios. The wealth is assumed to grow with the expected return, and follows the optimal control each period.

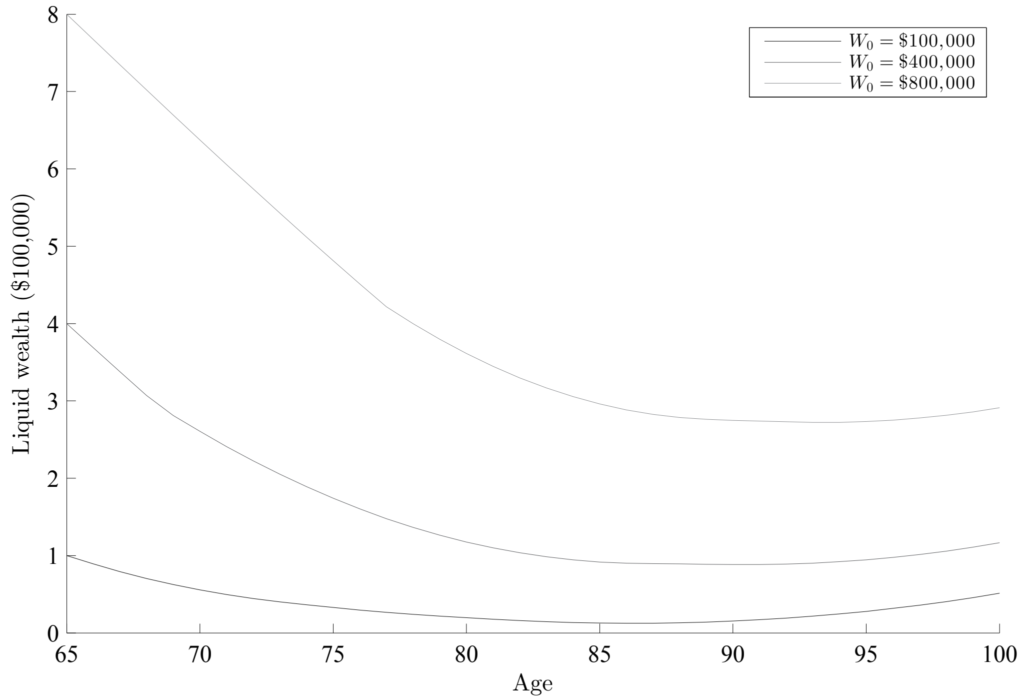

Figure 3.3: Wealth evolution for a single non-homeowner household given different starting wealth at $t = 65$, where wealth is drawn down based on optimal drawdown and grows with the expected risky return.

3.5 Conclusion

In this chapter, a sequential expected utility model was developed for the decumulation phase of Australian retirees. The model can be considered a more realistic extension to the work of Ding (2014), where stochastic wealth, stochastic family status, and a "health" status proxy were introduced, and prior assumptions of the order of means-test phases were relaxed. The model was defined as a stochastic control problem and solved for optimal consumption, optimal risky asset allocation, and optimal housing. The problem was solved numerically using dynamic programming.

The model explains the general behaviour of Australian retirees (as explained in Asher et al. (2017)), including declining consumption due to age, increased risk aversion with wealth, asset decumulation early in retirement, and asset accumulation later in retirement. The model can suggest an optimal policy for consumption and risky asset allocation with respect to the means-tested Age Pension.

The model lends itself to applications on both macro and micro scales. Once calibrated, it can be used to forecast future Age Pension needs and the implications of policy changes (or new financial products) on retirement behaviour. Moreover, it can be used on a financial planning level once individual risk preferences have been estimated. Further, it can easily be extended to suit the defined contribution pension system in other countries, or in the case of future changes to the Australian Age Pension.

Chapter 4

Calibration and analysis of Australian retirement behaviour

4.1 Introduction

Means-tested pension policies have become more important globally as the general population ages and life expectancy improves. The models used by the Australian Government to aid with planning and policy change tend to focus on a macro level (KPMG, 2010), and to evaluate the effect a policy change has on the individual retiree the policy makers must rely on empirical data after the policy has already been implemented. This chapter demonstrates how the assessment of policy changes can be done via an expected utility model in the case of the Australian system. The motivation for this research was the recent changes for Allocated Pension accounts, where assets in account-based pensions now generate a deemed income and no longer have an income-test deduction. Account-based pensions (such as Allocated Pension accounts) are accounts that have been purchased with Superannuation and generate an income stream throughout retirement. Prior to 2015, these types of accounts allowed for an income-test deduction that was determined upon account opening, and withdrawals were considered to be income in the means-test. The income-test deduction allowed the retiree to withdraw slightly more every year without missing out on the Age Pension. However, in 2015 the rules changed. Existing accounts were 'grandfathered' and will continue to be assessed under the old rules, while the new rules will be applied to any new accounts. The argument for the changes were simplicity (people with the same level of assets should be treated the same no matter how those assets are invested), to increase the incentive to maximise total disposable income rather than maximising Age Pension payments, and to simplify how capital growth and interest paying investments were assessed (Department of Social Services, 2016). From a fiscal point of view, the recommendations to introduce the new rules were based on estimated unchanged costs[1] (Henry, 2009). However the 2015–2016 budget stated expected savings of \$57m for 2015–2016, and \$129m and \$136m for subsequent years (The Commonwealth of Australia, 2015). The allocation to Age Pension in the 2015–2016 budget includes all changes to the Age Pension in a combined viewpoint, so a specific impact of the deeming rule changes is not known.

Another point of interest is the minimum withdrawal rules. They are designed to exhaust the retiree's account around year 100, however after year 85 (subject to investment returns) the withdrawn dollar amount starts decreasing quickly. In a recent report from Plan For Life (2016) it is identified that only 5% of retirees exhaust their accounts completely, though this number is expected to increase as life expectancy increases and the population ages. The report finds that retirees tend to follow the minimum withdrawal rules as guidelines for their own withdrawal, as few withdraw more than the minimum amount. This is further confirmed in Shevchenko (2016). Even so, Rice Warner (2015) argues that the minimum withdrawal rates should be cut by 25-50% to prevent retirees from exhausting their Superannuation prematurely due to increased longevity. The current rates are simply too high for many retirees, thus is not sustainable for people living longer than the average life expectancy, and are significantly higher than what is optimal in Andreasson et al. (2017).

There is very limited research modelling the Australian Age Pension, and even less research is conducted on implications of the regulatory minimum withdrawal rates even though a large number of retirees are using such accounts (or similar phased withdrawal products). The exception is Bateman and Thorp (2008), who

[1]The recommendations to introduce deeming was made in Henry (2009) where the fiscal sustainability is evaluated with the general equilibrium model 'KPMG Econtech MM900' (KPMG, 2010). The model shows the estimation over a 10-year window hence we do not know the short term or year-to-year estimates. In addition to this, the model includes additional suggested tax and budget related changes, hence the effect of introducing deeming rates cannot be isolated.

compare the welfare of retirees when the current minimum withdrawal rates were introduced in 2007 against the previous rules and alternative drawdown strategies. The authors use a rather simple CRRA model to examine the effect of different risk aversion and investment strategies but find that the minimum withdrawal rules increase the welfare for retirees although slightly less than optimal drawdown does. In Andreasson et al. (2017) the minimum withdrawal rules are included in part of the model outcome, but the analysis is by no means exhaustive and only provides a brief introduction to the effects. Ding (2014) does not constrain drawdown with minimum withdrawal, which would limit the author from finding a closed form solution. Other authors that focus on means-tested pensions do not enforce minimum withdrawal rates (Hulley et al., 2013; Iskhakov et al., 2015). It should be noted that their assumptions do not include Allocated Pension accounts, thus minimum withdrawal rates may not apply.

In order to evaluate policy changes, the model needs to be parameterised via calibration. Most research studies related to retirement models do not involve calibration against empirical data, which is crucial in order to use such models for conclusions or predictions in retirement. Tran and Woodland (2014) discussed the selection of model parameters that either match Australian economy rates or are taken from other studies to ensure that a general equilibrium steady state is achieved; however, the parameters are not calibrated against data. Calibration in Ding (2014) is based on the Mean-Squared Error, but the residuals are normalised with the estimated lifetime wealth. This allows for larger errors for poor households (but improved fit for wealthier households), as the estimated pension received will be rather large in relation to the actual wealth. Such normalisation may affect the outcome of the model, especially in forecasting future Age Pension budgets, as the error tend to be larger among the group that are the most dependent on the Age Pension. Ameriks et al. (2011) calibrated a bequest model using Maximum Likelihood Estimation to investigate the bequest motives of the US data by separating precautionary savings from bequests.

The chapter is structured as follows. Section 4.2 describes the data, assumptions and external parameters imposed for the calibration. Section 4.3 contains the statistical model used and describes the calibration procedure. Section 4.4 discusses the results and sensitivity of the calibrated parameters. General results and the differences in optimal decisions for recent Age Pension policies are then examined in Section 4.5. Finally, Section 4.6 presents the concluding remarks.

4.2 Calibration framework

The model is calibrated using a similar approach and the same data as Ding (2014), but with maximum likelihood estimation instead of minimising the mean-squared errors in order to estimate the utility parameters. The model used was presented in Chapter 3, and allows for three optimal decisions: consumption, risky asset allocation and housing allocation. However, the optimal risky asset allocation is not calibrated as the dataset lacks such information.

4.2.1 Dataset

For the dataset, the Household Expenditure Survey (HES) 2009–2010 and the Survey of Income and Household (SIH) 2009–2010 from the Australian Bureau of Statistics (Australian Bureau of Statistics, 2011) were used. This dataset has the limitation that data are only collected for private households, hence, retirees within assisted care facilities are excluded. Out-of-pocket health expenses are included, but any costs associated with private facilities are not. It does not provide a specific age either, as retirees are grouped into five year groups (65–69,70–74,75–79) and all retirees 80 years or older are grouped together. It is therefore not possible to distinguish the change in behaviour during the last 20 years of the model. Furthermore, households with or without dependants are treated as one group, even if the consumption might differ in reality.

The households are filtered by labour status ('not in workforce') and the age requirement for the Age Pension to find eligible retirees, where the age of a couple household is based on the youngest spouse. The data are then aggregated for each household to reflect the total expenditure (excluding mortgage payments), family home value, and wealth. In order to eliminate possible reporting errors from the data, samples that received no Age Pension despite being entitled to a material portion are removed, as well as any samples with expenditure either less than $3,000 per year or greater than assets available (liquid assets and Age Pension). Since the model is restricted to no borrowing, any entries with negative wealth are removed as well. Thus, 2,017 samples are obtained for couple households and 2,038 samples for single households. Further details regarding how the data were aggregated can be found in Appendix A.

4.2.2 Assumptions

In order to ensure a realistic calibration, some assumptions and constraints are needed. The following assumptions are imposed:

- For samples where the age is above the entitlement age, it is assumed that the wealth at retirement was the same as the current wealth of the sample, in order to calculate the income-test deduction $M(t)$. This is done because the value of the income-test deduction is determined when the Allocated Pension account is opened, which is assumed to be at retirement.

- It is assumed that households are aware of their life expectancy and can hence take this into consideration for decisions.

- As the first wealth quartile in the dataset is unlikely to constitute homeowners, a lower threshold for housing is set to $30,000 to make this consistent with the data. A retiree with wealth below this level can therefore not be a homeowner.

The following constraints are imposed on the utility parameters to ensure meaningful variables rather than over-fitting the model to the sample data. Nevertheless, none of the constraints were binding for the parameters once the calibration was completed.

$\bar{c}_d \in [0, P_{\max}]$ ensures that the consumption floor does not violate budget constraints for poor households. In other words, the necessary spending cannot be higher than income in the case of no accumulated wealth.

$\gamma_d < 0$, since the utility function is discontinuous at $\gamma_d = 0$, and positive values indicate risk-seeking behaviour.

$\theta \geq 0$, the preference of bequest over consumption must be positive as we do not allow negative bequest.

$a \geq 0$, the threshold for luxury bequest cannot be negative.

$\lambda_d \geq 0$, the utility parameter for being a homeowner cannot be negative; otherwise, the optimisation might suggest that selling a house that the retiree does not own while still receiving utility is optimal.

Minimum withdrawal rates and re-investments

There are minimum regulatory withdrawal rates imposed on Allocated Pension accounts, which increase with age[2]. It can be argued that this should be enforced in the calibration; however, it is intentionally left out for the following reasons. First, forced withdrawals affect whether the consumption consists of drawdown of wealth or Age Pension received, but not necessarily the level of consumption. Second, in response to the global financial crisis of 2007–2008, the government provided pension drawdown relief by reducing the minimum withdrawal rates between 2008 and 2013. As the dataset is taken from 2009–2010, the standard minimum withdrawal rates cannot be enforced during calibration. Finally, one advantage of an Allocated Pension account (the income-test deduction) is no longer available for new accounts after the 1st of January, 2015; hence, new retirees might opt for a different kind of account. To avoid limiting the analysis to pre-2015 account-based pensions, it was opted not to enforce minimum withdrawal rates in the model.

In addition, a retiree cannot deposit funds into an Allocated Pension account after retirement, where earnings are not taxed. However, the model allows the retiree to save Age Pension payments that have not been consumed, which would need to be invested in a separate account that might be subject to income-tax on returns. This would only occur for retirees with a very small amount of wealth; hence, the tax will be insignificant in the optimal decisions and not included in the model.

[2]The minimum withdrawal rate starts at 4% at age 65 and ends at 15% at age 95 (https://www.ato.gov.au).

Survival probabilities

The sequential model introduces two extra dimensions to the calibration: the age of the second person in a couple household, and the age difference between the spouses. In addition, survival probabilities differ between females and males, which would require additional parallel solutions to the model. If the age difference in couples was considered, the calibration would become too computationally expensive and unrealistic. Therefore, survival probabilities are generalised for single and couple households into a single unisex dimension.

To estimate the unisex survival probability, the ratio of males to females alive (estimated from the cumulative probability of being alive) is used as a weight. The estimated ratios match the empirical data proportions of males and females alive at any age t in Australian Bureau of Statistics (2011), and they can be used as a proxy for survival probabilities to avoid the gender variable. The unisex probability of surviving for one more year at age t is defined as

$$p_t^{\mathrm{S}} = 1 - \frac{q_t^{\mathrm{M}} \times {}_tp_0^{\mathrm{M}} + q_t^{\mathrm{F}} \times {}_tp_0^{\mathrm{F}}}{{}_tp_0^{\mathrm{M}} + {}_tp_0^{\mathrm{F}}}, \tag{4.1}$$

where the superscript indicates a single unisex household (S), male (M), or female (F) probabilities, q_t is the probability of dying before $t+1$ at age t and ${}_tp_0$ is the probability of surviving from birth (year 0) to age t. The actual mortality probabilities are taken from the Life Tables published in Australian Bureau of Statistics (2012).

The assumptions for couple households are different because it is already known that both spouses are alive; hence, no weighting of male-to-female ratios is necessary. The events of independent deaths of the spouses are non-mutually exclusive; however, they are treated as mutually exclusive to follow the model assumptions. We do not expect both spouses to die in the same year owing to the low probability of this occurring, and the effect it would have on the solution is minimal. The probability of surviving for one more year as a couple at time t is therefore defined as

$$p_t^{\mathrm{C}} = 1 - \left(q_t^{\mathrm{M}} + q_t^{\mathrm{F}}\right). \tag{4.2}$$

Portfolio composition and returns

The expected return and volatility have an important effect on the model calibration. To make the wealth process as realistic as possible, a typical portfolio composition of Self-Managed Super Fund (SMSF) accounts is estimated and then used with longer-term financial data to find the portfolio returns. Thus, the actual portfolio returns can be used in the calibration rather than returns based on the optimal allocation control parameter (which most likely will not be a correct representation for the average retiree).

To estimate the typical portfolio composition, SMSF data are used for each financial year from 2008 to 2014[3], from which actual investment returns on SMSF accounts are calculated. The average operating expense ratio reported by the ATO for the same period as the SMSF data is 0.83%[4]. The portfolio is assumed to be based on a set of risky assets approximated with the S&P/ASX 200 Total Return, which includes dividends, and a risk-free asset approximated with the deposit interest rate. The portfolio weight (proportion of risky assets) δ is then estimated with least squares regression. The average SMSF account returns for each year are regressed on the returns of the S&P/ASX 200 Total Return and the deposit rate. The risky allocation δ is found to be 43.7% with a significance level of 1%, which is used as a proxy for risky asset allocation during calibration of the model.

The long-term return estimate is taken as the 20-year average log-returns prior to 2010 of the S&P/ASX 200 Total Return. The returns are then adjusted to real returns by deducting inflation ($\widetilde{r} = 2.9\%$) and for operating expenses. The final estimates are $r_t = 0.005$ and $Z_t \sim \mathcal{N}(0.056, 0.018)$, which are used in equation (3.1).

4.2.3 Age Pension and parameters

The parameters for the Age Pension are listed in Table 4.1 and taken for the year 2010 to match the data. A retiree is eligible for the Age Pension at age 65 (male) or 63 (female). The scaling factor for couple households is set to $\zeta = 1.3$, which is in line

[3]Data contain individual account balances, earnings, and drawdowns each year for Self-Managed Super Fund accounts, taken from a dataset provided by the Australian Taxation Office to CSIRO-Monash Superannuation Research Cluster (not publicly available).

[4]Data are taken from 'Australian Taxation Office Reports SMSFs: A statistical overview' for the years 2008–2009, 2009–2010, 2010–2011, 2011–2012, and 2012–2013.

with the results of Fernández-Villaverde and Krueger (2007), who reviewed research on controlling for family size and the resulting economy of scale. In addition, the terminal age is set to $T = 100$. As the Allocated Pension account was operating under the policy rules that allowed for an income-test deduction, and where income was considered equal to the amount drawn down from the account, the formula for Age Pension defined in equation (3.17) is used with the income-test definition in equation (3.13).

Table 4.1: Age Pension rates published by Centrelink as at January 2010 (*www.humanservices.gov.au/customer/services/centrelink/age-pension*).

		Single	Couple
$P^d_{\max}$	Full Age Pension per annum	\$17,456	\$26,099
	Income-Test		
L^d_{I}	Threshold	\$3,692	\$6,448
ϖ^d_{I}	Rate of Reduction	\$0.5	\$0.5
	Asset-Test		
$L^{d,h=1}_{\mathrm{I}}$	Threshold: Homeowners	\$178,000	\$252,500
$L^{d,h=0}_{\mathrm{I}}$	Threshold: Non-homeowners	\$307,000	\$381,500
ϖ^d_{A}	Rate of Reduction	\$0.039	\$0.039

4.3 Calibration model and procedure

Calibration of the model with the sample data are performed with the maximum likelihood method. The sample data are split into vectors of single ($d = \mathrm{S}$) and couple ($d = \mathrm{C}$) households for consumption ($\mathbf{c}^d$) and housing ($\mathbf{h}^d$). Denote the total data as $\mathcal{D} = \{\mathbf{c}^{\mathrm{S}}, \mathbf{c}^{\mathrm{C}}, \mathbf{h}^{\mathrm{S}}, \mathbf{h}^{\mathrm{C}}\}$. The statistical models are assumed to be $c^d_i = \widetilde{c}^d_i(\Theta)e^{\epsilon^d_i}$ and $h^d_i = \widetilde{h}^d_i(\Theta)e^{\varepsilon^d_i}$, where c^d_i and h^d_i are sample i from the corresponding dataset, $\widetilde{c}^d_i(\Theta)$ and $\widetilde{h}^d_i(\Theta)$ are the optimal consumption and housing, respectively, from the model based on wealth and age corresponding to sample i, and the utility model parameter vector is

$$\Theta^{\top} = \begin{pmatrix} \gamma_{\mathrm{S}} & \gamma_{\mathrm{C}} & \gamma_{\mathrm{H}} & \theta & a & \bar{c}_{\mathrm{S}} & \bar{c}_{\mathrm{C}} & \psi & \lambda \end{pmatrix}. \tag{4.3}$$

Finally, $\epsilon^d_i \sim \mathcal{N}(0, (\sigma^d_\epsilon)^2)$ and $\varepsilon^d_i \sim \mathcal{N}(0, (\sigma^d_\varepsilon)^2)$ are the independent non-standardised error terms (residuals).

The log likelihood function of N independent identically distributed samples $\mathbf{x} = (x_1, x_2, ..., x_N)$ from a log-normal distribution such that $\ln x_i \sim \mathcal{N}(\ln \mu, \sigma^2)$ is

$$\mathcal{L}_{\mathbf{x}}(\mu, \sigma) \propto -\frac{N}{2} \ln\left(2\pi\sigma^2\right) - \frac{1}{2\sigma^2} \sum_{i=1}^{N} (\ln x_i - \ln \mu)^2 . \tag{4.4}$$

The maximum likelihood estimates of parameters Θ and $\boldsymbol{\sigma} = \{\sigma_\epsilon^{\mathrm{S}}, \sigma_\epsilon^{\mathrm{C}}, \sigma_\varepsilon^{\mathrm{S}}, \sigma_\varepsilon^{\mathrm{C}}\}$ are obtained by maximising the total log likelihood:

$$\mathcal{L}_{\mathbf{c}^{\mathrm{S}}}(\widetilde{\mathbf{c}}^{\mathrm{S}}(\Theta), \sigma_\epsilon^{\mathrm{S}}) + \mathcal{L}_{\mathbf{c}^{\mathrm{C}}}(\widetilde{\mathbf{c}}^{\mathrm{C}}(\Theta), \sigma_\epsilon^{\mathrm{C}}) + \mathcal{L}_{\mathbf{h}^{\mathrm{S}}}(\widetilde{\mathbf{h}}^{\mathrm{S}}(\Theta), \sigma_\varepsilon^{\mathrm{S}}) + \mathcal{L}_{\mathbf{h}^{\mathrm{C}}}(\widetilde{\mathbf{h}}^{\mathrm{C}}(\Theta), \sigma_\varepsilon^{\mathrm{C}}), \tag{4.5}$$

with respect to $(\Theta, \boldsymbol{\sigma})$.

The calibration is carried out in two steps. First, a suitable starting point is identified by searching globally. Each utility parameter is assigned a realistic range of values where three different values are selected. Once a starting point is identified, the parameters are optimised further using the Nelder-Mead simplex algorithm until further improvements of the log likelihood function are negligible.

The calibration is computationally expensive. Each time the parameters are updated the model needs to be solved four times (each combination of single/couple households and homeowner/non-homeowner). Then, the model output needs to be generated from the wealth, age, and homeowner data given in each sample. The model output is compared with the sample data for consumption and housing to estimate the fit of the model before the next iteration begins.

4.4 Calibration results

The calibration output indicates that the model fits the empirical behaviour. The statistical model is well chosen with respect for consumption, but housing residuals do not follow a normal distribution. Figure 4.1 shows a Quantile-Quantile plot of the residuals for couple households who are homeowners.

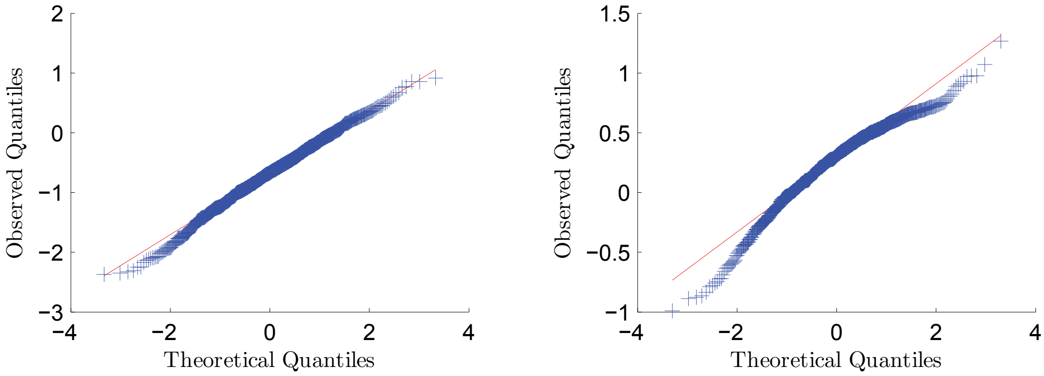

(a) Consumption residuals (b) Housing residuals

Figure 4.1: Quantile-Quantile plot for couple households where the residuals are assumed to follow a normal distribution.

Other distributions, such as skew-t, lead to improvement in residual fitting (Figure 4.2). The residuals are then assumed to be $\epsilon_i^d \sim \mathcal{ST}_\nu(0, \sigma_\epsilon^d, \omega)$ and $\varepsilon_i^d \sim \mathcal{ST}_\nu(0, \sigma_\varepsilon^d, \omega)$ with ν degrees of freedom and skewness ω, which are estimated from the residuals. This improves the fit of housing residuals, although extreme values (larger than four standard deviations) will result in failing normality tests. The assumption of normally distributed residuals generates $AIC = 4104$, while skew-t residuals $AIC = 682$. The assumption that consumption and housing residuals are independent is confirmed.

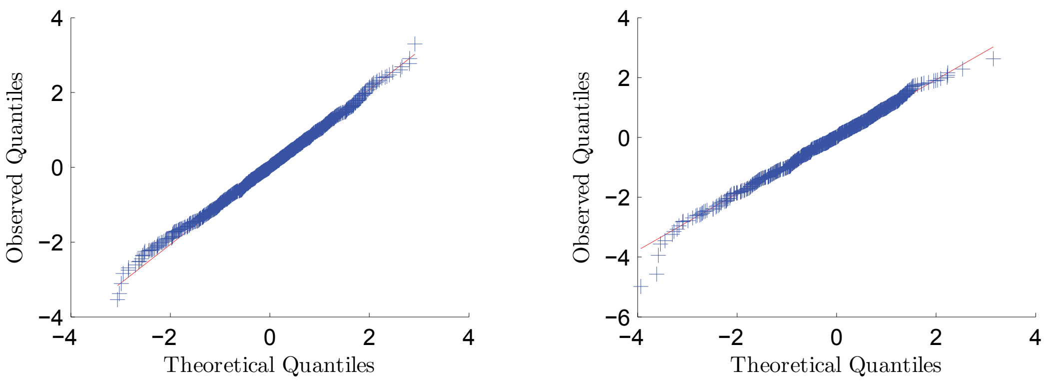

(a) Consumption residuals (b) Housing residuals

Figure 4.2: Quantile-Quantile plot for couple households where the residuals are assumed to follow a skew-t distribution.

4.4.1 Calibrated parameters

The estimated parameters (see Table 4.2) are in line with the related literature. The risk aversion is slightly lower than that of Ding (2014), who estimated $\gamma = -3$, while the consumption floor is well below the full Age Pension rates but in line with the author's findings. In relation to Asher et al. (2017), the consumption floor is marginally lower than the lower quintile in 'Average Household Consumption' for singles (\$11,900) and for couples (\$21,400). As the estimate in Asher et al. (2017) is not actually a consumption floor (even the lowest quintile will have some assets), the model output for total consumption will be similar to these estimates. This can be compared with Ameriks et al. (2011), where the consumption floor is only \$5,750 USD (the average US social security payment): hence, in relative terms, the calibrated estimates are higher.

Table 4.2: Calibrated parameters with standard errors.

	γ_S	γ_C	γ_H	θ	a	$\bar{c}_S$	$\bar{c}_C$	ψ	λ
Value	-2.77	- 2.29	-2.58	0.54	26 741	11 125	18 970	1.47	0.037
Std. Error	0.12	0.14	0.19	0.03	1 377	1 011	1 682	0.04	0.006

The calibrated "health" proxy parameter $\psi = 1.47$ indicates that the preferences for consumption that exceed the consumption floor decrease with a factor of $\psi^{1/(\gamma_d - 1)}$; hence, a 10% decrease for singles and 11% for couples each year owing to declining health. To put this into perspective, it equals a decrease of \$1,828 per year for the median single household at age 65, and \$3,198 for the median couple household. Bernicke (2005) conducted a US-based study on empirical retirement data and found that the difference in consumption for the age span 65-75 was 26.4%, and Higgins and Roberts (2011) found a 20%-30% drop in median levels where the decrease is larger for wealthier households, which confirms the results. Clare (2014) found an average yearly decrease of 10% for households with a comfortable lifestyle (roughly 300% higher than the calibrated consumption floor) and 2% for a modest lifestyle (roughly 100% higher than the calibrated consumption floor) between ages 70 and 90. The calibrated model suggests a similar decrease in expenditure and captures the characteristics of larger declines for wealthier households.

Note that the housing preference (λ) no longer have subscripts that separate between single (λ_S) and couple (λ_C) households. During calibration the two alternative parameters ended up moving together, and the results showed the same value down to three decimals, thus are now treated as a single parameter.

Finally, θ implies medium sensitivity of consumption to wealth, which implies that additional wealth is being saved for a bequest. The parameter estimate is different from Ding (2014) ($\theta = 0.956$), where the calibrated parameter is lower due to housing now being included in bequest. It indicates that a sequential model does not affect bequest motives, but adjusts the preferences to account for the increase in expected bequeathed asset due to housing. Nevertheless, the model cannot distinguish whether this is due to precautionary savings or indeed clear bequest motives.

4.4.2 Parameter sensitivity

In order to estimate the sensitivity of the calibrated utility parameters the impact on the optimal decision parameters is measured. Due to difficulties in calculating derivatives, the analysis relies on numerical perturbation where each parameter is adjusted with ± 2 standard errors from Table 4.2, and the largest absolute change from a benchmark is shown for each decision variable in Table 4.3. As a benchmark, a couple household is chosen at the time of retirement, with a total initial wealth of \$1m. The optimal decisions for the benchmark parameters equals $\alpha_{t_0} = 0.17$, $\delta_{t_0} = 0.72$ and $H = \$554,744$.

The changes in risk aversion for couple households (γ_C) and housing (γ_H) tend to have the largest impact, while the changes in risk aversion for single households (γ_S) tend to be negligible. When using two standard errors this places γ_H outside the other risk aversion parameters, hence marginal utility for housing increase or decrease significantly and affects housing allocation but not the other decision variables. Similarly, changes in γ_C either increases or decreases the distance relative to the other risk aversion parameters, hence relative preferences to housing changes a lot, although risk aversion with regard to risky asset investments also changes notably. The consumption floors, "health" proxy and housing preferences show less sensitivity, and the bequest parameters close to none.

Note that the solution of the utility model depends on the relative marginal utility between functions and over time, hence similar changes to all risk aversion parameters (e.g. if all three were adjusted with +2 standard errors) the sensitivity would be substantially lower and optimal decisions would change less. It is our opinion that the model therefore is fairly robust with regards to the parameterisation.

Table 4.3: Sensitivity of control variables when calibrated parameters are adjusted ± 2 standard errors.

Parameter	$\max\lvert\Delta_\alpha\rvert$	$\max\lvert\Delta_\delta\rvert$	$\max\lvert\Delta_H\rvert$
γ_{S}	0.00	0.00	\$1,309
γ_{C}	0.03	0.19	\$246,610
γ_{H}	0.01	0.00	\$645,256
$\bar{c}_{\text{S}}$	0.00	0.09	\$1,106
$\bar{c}_{\text{C}}$	0.02	0.00	\$95,653
a	0.00	0.00	\$1
θ	0.00	0.02	\$1
ψ	0.01	0.05	\$9,278
λ	0.01	0.07	\$96,119

4.4.3 Shortcomings of calibration

We acknowledge the disadvantages of fitting a multi-period model on cross-sectional data, as cohort effects are not accounted for. Because the Australian Age Pension system is yet to mature, it is likely that the wealth for each age group is underestimated compared to a mature system, especially at an older age. Although it is difficult to anticipate the effect this would have on the calibration, it is reasonable to expect that if consumption remains near current levels, then the risk aversion parameter (γ) would increase, while if wealth is underestimated only for older retirees, then the "health" proxy (ψ) would increase. This would mean lower consumption in relation to wealth, as well as an overall lower proportion of risky asset allocation. In addition, the calibration suffers from 'identification problems': different utility parameters can result in similar outputs as it is the ratio of marginal utility between utility functions that affects optimal decisions, hence the actual value of the utility function has a less significant effect. This is again due to the data being a cross-sectional dataset, hence cohort effects cannot be distinguished.

4.5 Analysis of Age Pension policy

The calibrated parameters are now used to compare the impact that four different means-tested policies have on an Australian retiree. Owing to the limitations in the data used for calibration (see Section 4.2.1), the result should only be considered for healthy households in the post-retirement phase. This is because the survey samples tend to be of better health than what might be true in reality, as retirees with poor health are more likely to live in an assisted care facility.

As the means-test thresholds are adjusted over time, it is important to adjust the parameters accordingly to allow for different policies to be compared - especially since the calibration was carried out on data from 2010. The calibrated consumption floor $\overline{c}_d$ and the threshold for luxury bequest a must be adjusted as they represent monetary values. Since the model is defined in real terms, a new base year must be set for the comparison. Parameters are therefore adjusted based on the Age Pension adjustments from 2010 to 2017. Currently, the Age Pension payments are adjusted to the higher of the Consumer Price Index and Male Average Weekly Total Earnings. The increase in full Age Pension payments from 2010 to 2017 equals approximately a 4.5% increase per year. It is assumed that the utility parameters representing monetary values have increased in the same manner, hence adjusted to $\overline{c}_{\mathrm{S}} = \$13,284$, $\overline{c}_{\mathrm{C}} = \$20,607$ and $a = \$27,200$. Any other parameterisation remains the same as in the calibration.

4.5.1 Policy definitions

The focus is on four different policies that represent recent changes in the Australian Age Pension system with respect to model assumptions. A summary of the policies with the Age Pension rates and means-test assumption is shown in Table 4.4.

Policy 1 - Pre January 2015, no minimum withdrawal (PRE2015NOMWD)

The first policy reflects the means-test and policy rules that applied for the period of the calibration data. The actual thresholds in the means-test have been adjusted for the base year 2017. The income in the income-test is based on drawdowns from the Allocated Pension account, and no minimum withdrawal rules are enforced. Since

Table 4.4: Age Pension rates and rules used for policy variations.

	PRE2015NOMWD	PRE2015	POST2015	POST2017
$P^{\mathrm{S}}_{\max}$	\$22,721	\$22,721	\$22,721	\$22,721
$P^{\mathrm{C}}_{\max}$	\$34,252	\$34,252	\$34,252	\$34,252
Income-Test	Drawdown	Drawdown	Deemed	Deemed
L^{S}_{I}	\$4,264	\$4,264	\$4,264	\$4,264
L^{C}_{I}	\$7,592	\$7,592	\$7,592	\$7,592
ϖ^{d}_{I}	\$0.5	\$0.5	\$0.5	\$0.5
κ^{S}	-	-	\$49,200	\$49,200
κ^{C}	-	-	\$81,600	\$81,600
ς_{-}	-	-	1.75%	1.75%
ς_{+}	-	-	3.25%	3.25%
Asset-Test				
$L^{S,h=1}_{\mathrm{A}}$	\$209,000	\$209,000	\$209,000	\$250,000
$L^{C,h=1}_{\mathrm{A}}$	\$296,500	\$296,500	\$296,500	\$375,000
$L^{S,h=0}_{\mathrm{A}}$	\$360,500	\$360,500	\$360,500	\$450,000
$L^{C,h=0}_{\mathrm{A}}$	\$448,000	\$448,000	\$448,000	\$575,000
ϖ^{d}_{A}	\$0.039	\$0.039	\$0.039	\$0.078
Min. Withdrawal	No	Yes	Yes	Yes

minimum withdrawals were not enforced in the calibration, the optimal behaviour is examined without the requirement to withdraw a certain amount from the Allocated Pension account each year.

Policy 2 - Pre January 2015 (PRE2015)

The difference from Policy 1 is that minimum withdrawals are now enforced (and applied to policies 3 and 4 as well). The retiree can no longer opt to withdraw a lesser amount from the wealth in order to receive a full or partial Age Pension to cover consumption needs. Minimum withdrawal rates for Allocated Pension accounts are shown in Table 4.5 (Australian Taxation Office, 2016). The rates impose a lower bound on optimal consumption, hence withdrawals from liquid wealth must be larger or equal to these rates.

Policy 3 - Post January 2015 (POST2015)

The focus of the changes introduced on the 1$^{\text{st}}$ of January 2015 was on the income-test and Allocated Pension accounts. The income-test now uses deemed income rather than drawdown, thus the liquid wealth is used in both the asset and income-test. The retiree can therefore withdraw more liquid wealth without missing out on Age Pension payments.

Policy 4 - Asset-test changes January 2017 (POST2017)

On the 1^{st} of January 2017 the thresholds of the asset-test were 'rebalanced', hence changed significantly (Australian Government Department of Veterans' Affairs, 2016). The thresholds for the asset-test increased and the taper rate ϖ_{A}^{d} doubled. This effectively means that retirees will now receive a full Age Pension for a higher level of wealth, but once the asset-test binds, the partial Age Pension will decrease twice as fast, causing them to receive no Age Pension at a lower level of wealth than before. No adjustments were made to the full Age Pension or income-test threshold.

Table 4.5: Minimum regulatory withdrawal rates for Allocated Pension accounts for the year 2016 and onwards (*https://www.ato.gov.au/rates/key-superannuation-rates-and-thresholds/*).

Age	≤ 64	65-74	75-79	80-84	85-89	90-94	$95 \leq$
Min. drawdown	4%	5%	6%	7%	9%	11%	14%

4.5.2 Age Pension function

The Age Pension function was defined in 3.2.4. It can be seen in Figure 4.3 how each policy affects the Age Pension received, given a certain liquid wealth at age 65 for a single non-homeowner household. As the PRE2015 policy depends on the drawdown of wealth as income, the minimum drawdown level of 5% is used. The shape of the curves is very similar for couples and non-homeowners, although with different thresholds hence the kinks appear at different wealth levels. It should be noted that the similarities for POST2015 and POST2017 below ~\$550 000 are because the deemed income from assets makes the income-test binding in both policies. In the case of a non-financial asset, such as a boat or jewellery, then POST2017 would generate more Age Pension for lower levels of wealth than POST2015.

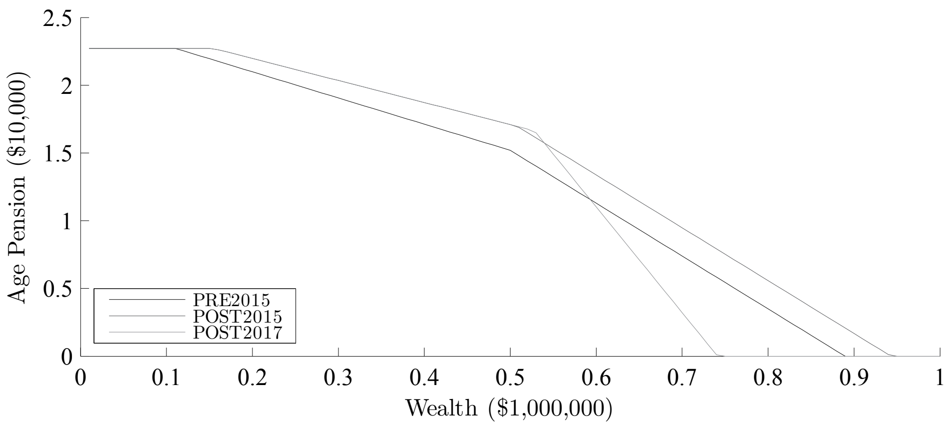

Figure 4.3: Comparison of Age Pension function under different policies for a single non-homeowner household aged 65.

4.5.3 Optimal consumption

The optimal consumption curve in relation to wealth differs from the one in traditional utility models, where the deviations can be explained by the Age Pension means-test parameters. Figure 4.4 shows optimal consumption in relation to liquid wealth for different age groups and Age Pension policies. The curves are very similar for both single households and couple households, hence only single households are shown. Traditionally, consumption is a smooth, concave, and monotone function of wealth and becomes flatter as wealth increases due to decreasing marginal utility. In general, this is true for drawdown outside the upper thresholds of the means-test where no Age Pension is received. As the means-test binds, indicated by the shaded background, the optimal behaviour changes slightly to anticipate the Age Pension received (which equals the area between the dashed and solid curves). This effect can be seen in the PRE2015NOWD and PRE2015 policies, due to drawdown of wealth being considered income for the means-test. For such policies it is optimal to withdraw slightly more when a retiree receives no Age Pension but is close to receiving partial, as larger withdrawals will decrease wealth (and in turn decrease the withdrawals) to make the retiree eligible to receive a partial Age Pension. For POST2015 and POST2017 policies the same effect does not exist. The asset-test is the first to bind as the retiree moves from no Age Pension to partial, so withdrawal decisions have less impact. Only the decision to withdraw less when receiving a partial Age Pension is present, which effectively replaces consumption from assets with consumption from the Age Pension. There is a marginal effect when the retiree goes from no Age Pension to receiving a partial Age Pension, especially for the 2017 asset-test adjustment, shown as a tiny dent where the consumption and drawdown curve intersect (the threshold between no Age Pension and a partial Age Pension due to the asset-test). This implies that a retiree should consume slightly more when the wealth is close to this threshold in order to receive a partial Age Pension, but the additional utility would be so small that it is negligible in planning. A retiree could therefore better plan the withdrawals in order to maximise Age Pension payments under the old policies, as sensitivity to the means-test decreased when the deemed income rules were introduced. The effect of an optimal consumption strategy becomes weaker as the minimum withdrawal rates cross over 'unconstrained' optimal withdrawals, and for the older retirees the optimal strategy is only to withdraw the minimum required rate, thus the Age Pension simply adds to the consumption rather than being included in desired consumption. This is in line with Bateman et al. (2007), which finds that welfare decreases slightly when minimum withdrawal rules are enforced over unconstrained optimal withdrawals, especially for higher levels of risk aversion.

An interesting outcome is when the consumption paths over a lifetime are compared between the different policies (Figure 4.5). Since the optimal drawdown rules are very similar between policies, and minimum withdrawal rates quickly binds, the consumption in turn follows the same pattern. However, due to the income-test changes after 2015, more Age Pension is received in POST2015 and POST2017 in relation to wealth and drawdown assuming the deeming rates stay constant. This is especially true early in retirement, where higher consumption otherwise would result in no Age Pension. One of the reasons for changing the policy was for the government to generate savings, but the deeming rules will not have the desired outcome on Allocated Pension accounts unless the deeming rates increase. Only if the minimum withdrawals are removed (or at least decreased), which in turn could lead to lower withdrawals for given wealth levels, could current rates lead to Age Pension payments being less under the more recent policies[5]. With the current rules

[5]It should be noted that the findings are for the account-based pension only, as other products which do not enforce the minimum withdrawal rates could incur additional savings for the government under the newer rules.

(POST2017) the retiree will receive a significant amount of Age Pension over the lifespan, unless a larger level of wealth has been accumulated prior to retirement.

Wealth paths throughout retirement, however, are almost identical except for PRE2015NOMWD as the lack of minimum withdrawal rules makes it possible for risky return to grow wealth. The difference between the minimum withdrawal enforced policies is solely in consumption from additional Age Pension payments.

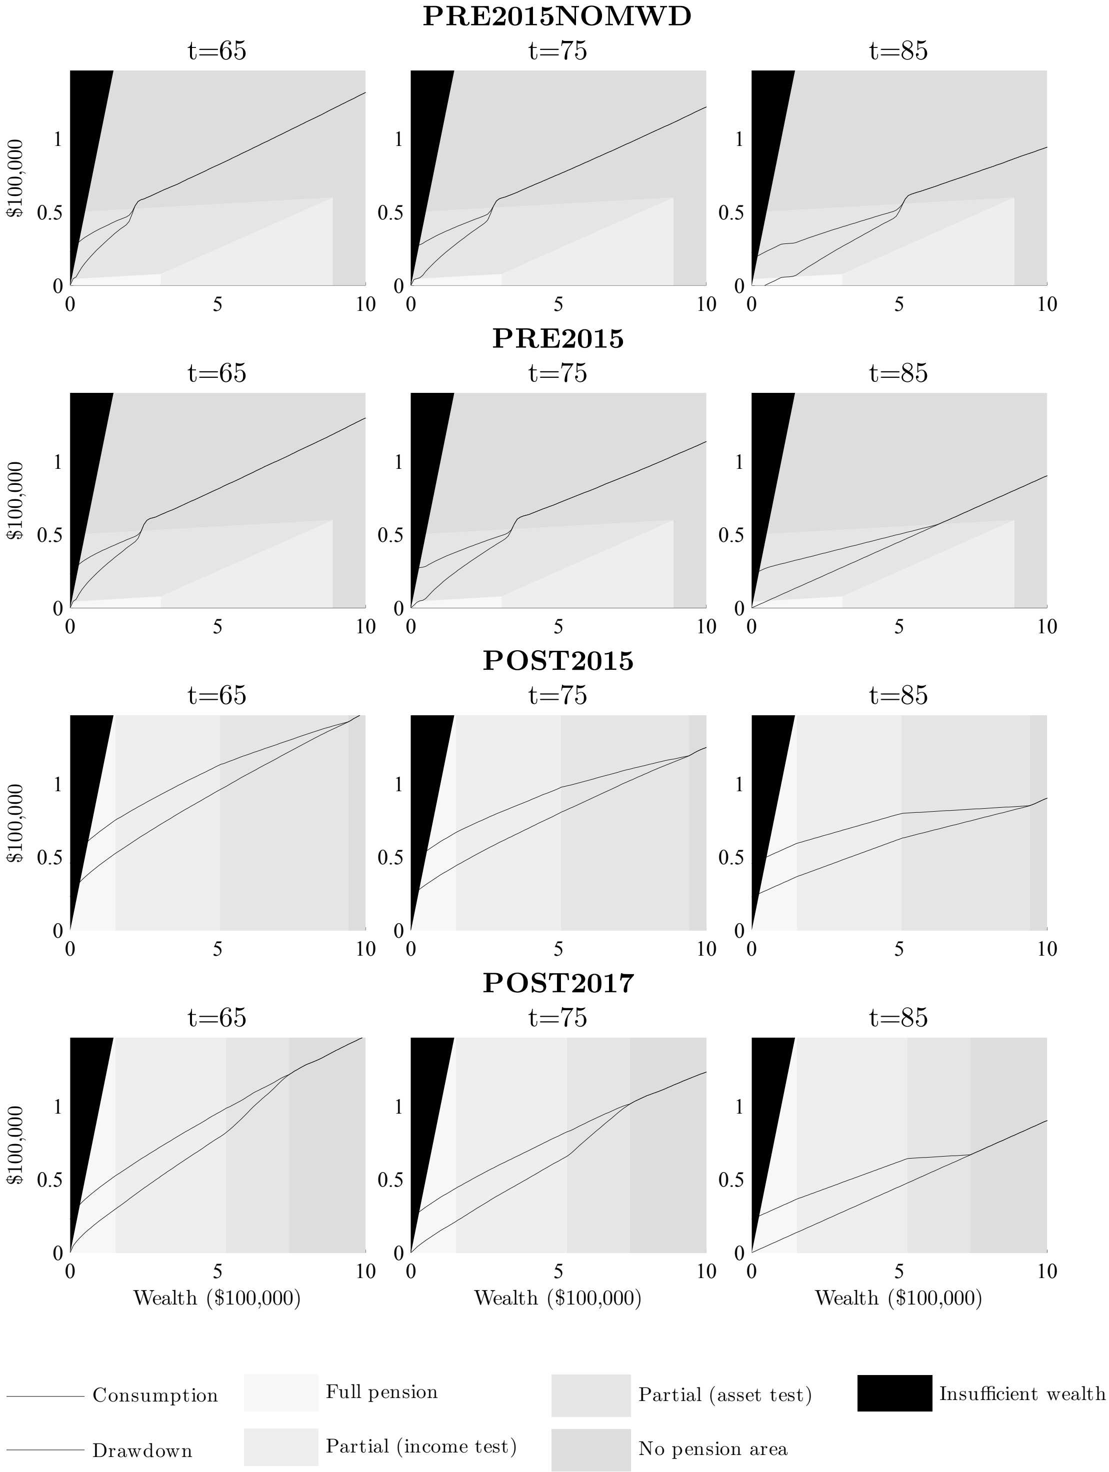

Figure 4.4: Optimal drawdown ($\alpha_t w_t$) and consumption in relation to liquid wealth for a single non-homeowner household, given different Age Pension policies and ages.

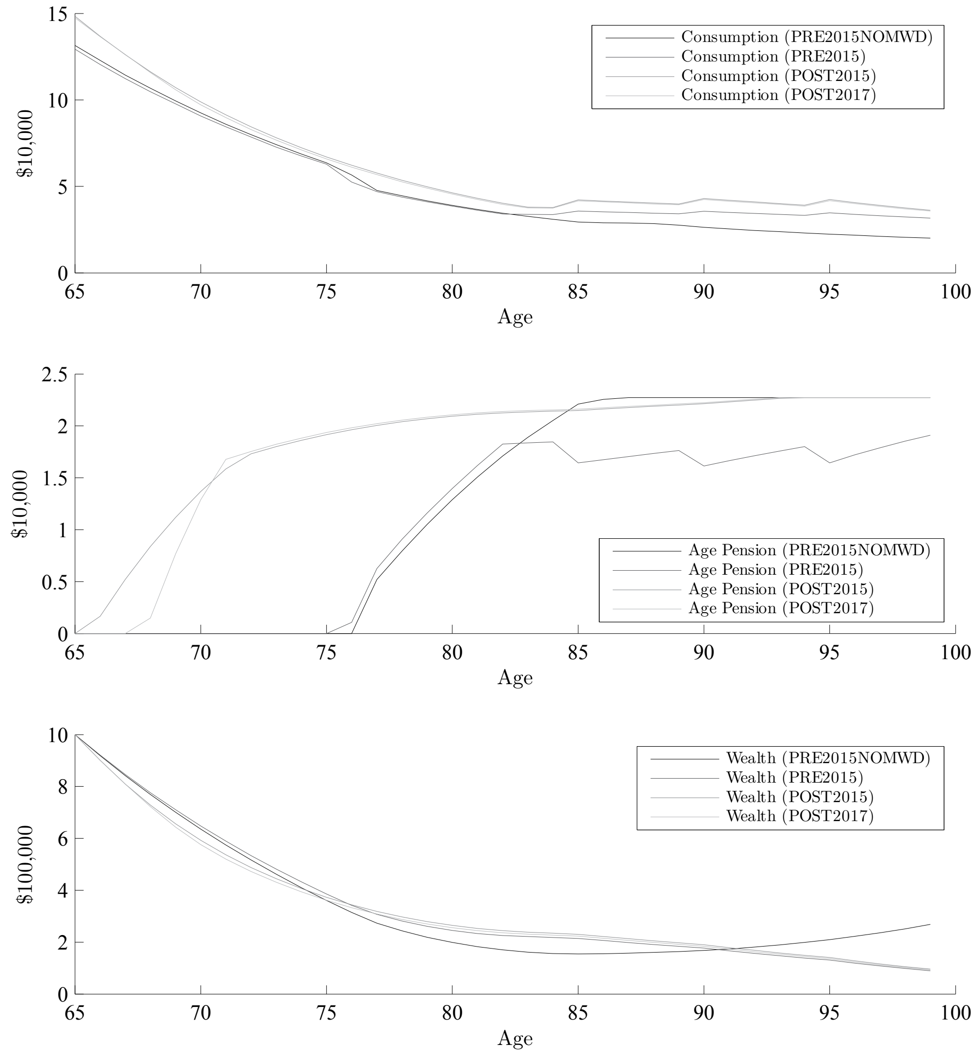

Figure 4.5: Comparison of consumption, Age Pension and wealth paths over a retiree's lifetime. The wealth grows with the expected return each year, and the drawdown follows the optimal drawdown decision each year.

4.5.4 Optimal risky asset allocation

Exposure to risky assets in the portfolio is highly dependent on wealth and age, and sensitive to the asset-test. In general, the percentage allocation tends to decrease with wealth and over time. Such behaviour agrees with traditional investment advice, which suggests that the allocation of risky assets should be reduced with age. However, it is contrary to the findings in Andreasson et al. (2017), where the same model was calibrated but where housing was not included as bequest, which was confirmed in Iskhakov et al. (2015) (as long as the consumption floor is less than the Age Pension) and Ding (2014). The latter showed that when bequest is considered a luxury, the optimal allocation of risky assets increases with age, implying higher allocation to risky assets throughout retirement. Decreasing exposure with age is indeed optimal only when bequest is not considered in a utility maximisation model (Blake et al., 2014). The reason why the calibration suggests this seemingly contrary behaviour comes from the fact that the preferences for risk are lower for a given liquid wealth in the bequest function than for consumption. As the mortality risk shifts the focus from consumption towards bequest, this moves the proportion risky allocation from high to low, therefore the optimal risky asset allocation decreases with age.

The contour charts in Figure 4.6 show a complex relationship with the means-test and Age Pension payments. The exposure to risky assets in the portfolio is highly dependent on wealth and age, and even more so for the more recent policies POST2015 and POST2017. This is expected since the means-test is now based on wealth in both the asset and the income-test, which means investment returns will have a larger impact on expected utility. The risky allocation can be explained with the expected marginal utility conditional on wealth. When marginal utility increases with wealth, the risky allocation will always suggest 100% risky assets. This is the case for the bottom black areas, where the upper bound to the left indicates the maximum marginal utility from consumption, and the upper bound to the right is the maximum marginal utility from bequest. If utility from consumption is considered individually, then lower levels of wealth will have higher marginal utility. If marginal utility from bequest is instead isolated, the same effect will occur albeit at a lower level than for consumption. It is, therefore, optimal up to these levels to allocate 100% to risky assets, as the potential return is valued more than the risk. The marginal utility is also affected by the means-test, as a result of the 'buffer' effect. This buffer occurs when the decreasing wealth that stems from an investment loss is partially offset via increased Age Pension payments and can be seen as the comparatively darker area towards the top left in POST2015 and POST2017. These areas correspond to just before a partial Age Pension is received. The buffer effect is, therefore, strongest for a retiree who has no Age Pension but is close to receiving a partial Age Pension. An investment loss in this instance would be offset by partial Age Pension, whereas an investment profit would not cause the retiree to miss out on Age Pension that he/she would otherwise receive. Hulley et al. (2013) found that risky allocation is much higher when the asset-test is binding owing to the steeper taper rate, and slightly lower (but still higher than the benchmark) for the income-test. This is the case especially for POST2017, hence, marginal utility is lower when the asset-test is binding. For very low levels of wealth, the buffer effect is the opposite; investment losses can never lead to more than full Age Pension, and investment profits will decrease the amount of partial Age Pension received, which will result in lower marginal utility. At this point the marginal utility from potential return outweighs the negative buffer effect, indicating full risky asset allocation being optimal.

In addition to the impact of the Age Pension, a few general conclusions can be derived with respect to optimal risky asset allocation. First, as wealth increases, the allocation to risky assets decreases. A loss in wealth has more negative marginal utility than the equivalent gain has positive utility; hence, preserving available wealth is preferred over potential additional wealth. This effect increases with wealth owing to the HARA utility. Second, as the retiree ages, the (mortality risk) weight increases towards bequest, where the marginal utility is lower. This implies that preserving capital becomes even more important with age as wealth increases, although the effect of being a homeowner skews this upwards as the home value effectively increases the threshold for luxury bequest. Finally, couples tend to be more aggressive with risky assets. One factor is their slightly lower risk aversion compared with singles, but because the risk of receiving bequest is lower for couples than for singles, couples are less affected by the second factor stated above. In addition, couples have a higher chance of recovering negative asset shocks; hence, they have more to gain from higher risk exposure in the long term.

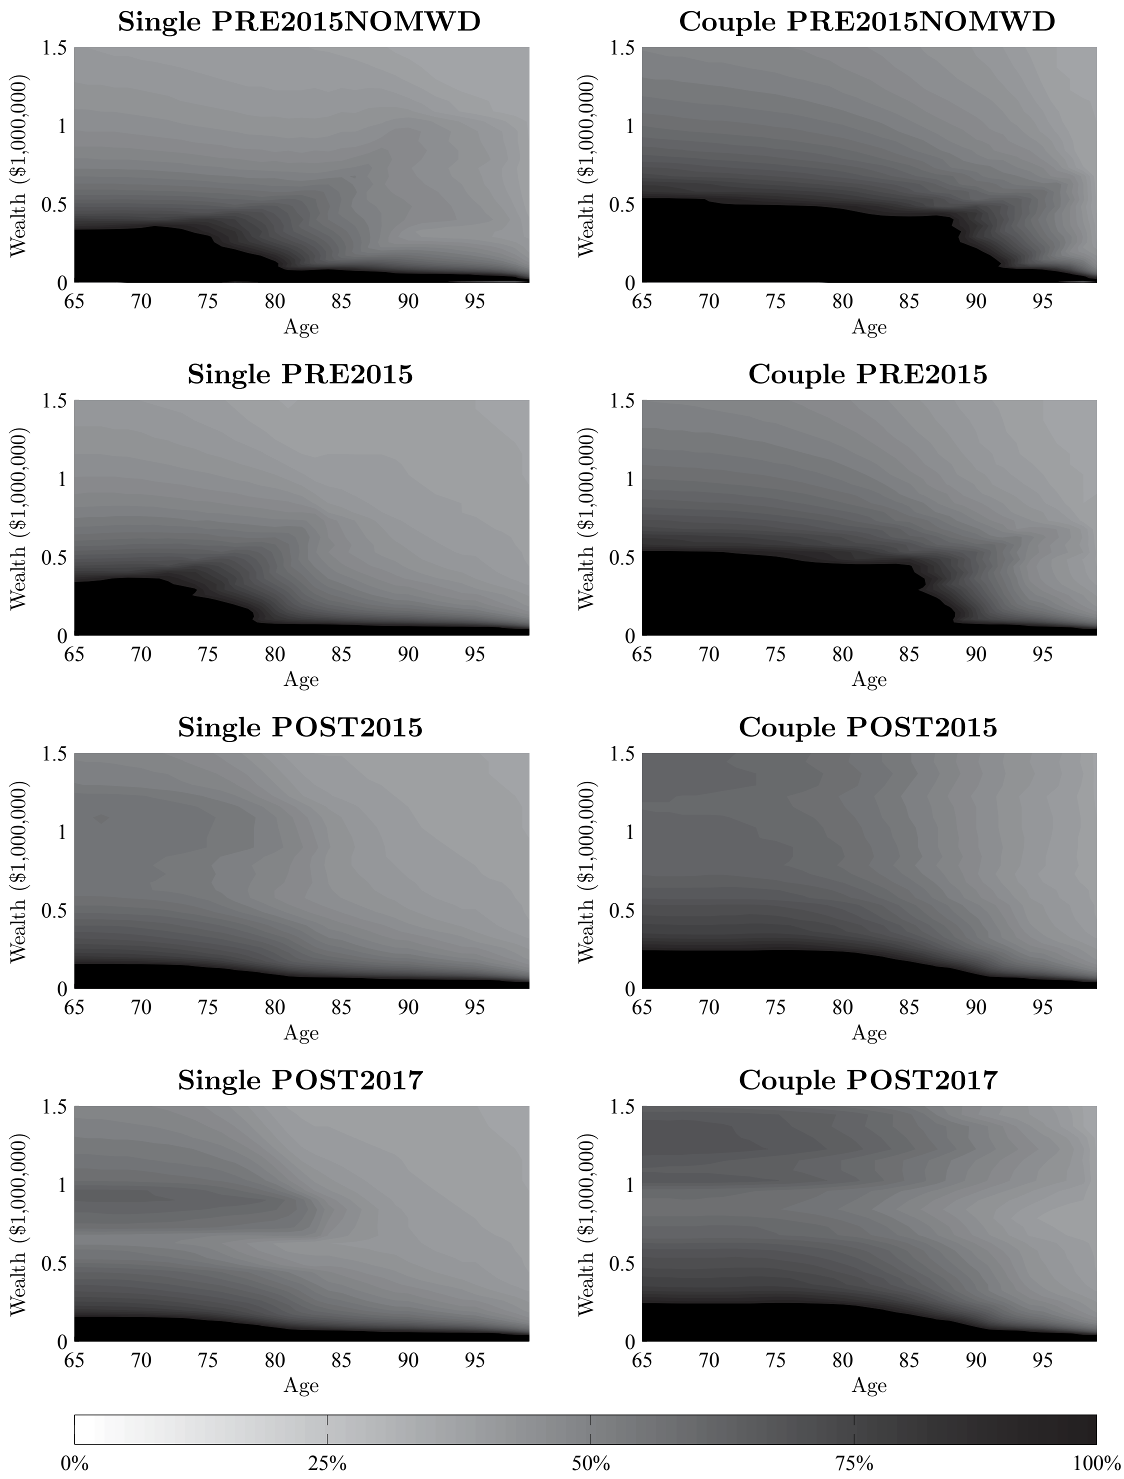

Figure 4.6: Optimal allocation to risky assets for single and couple non-homeowners given liquid wealth, for different Age Pension policies.

4.5.5 Optimal housing allocation

Optimal housing indicates that the retiree is better off allocating the majority of assets to housing, the reason being two-fold. First, the house value is not included in the means-test, hence the retiree can have significant assets and still receive Age Pension payments. Second, the house can be bequeathed, hence the retiree can consume all liquid assets and still receive utility by leaving bequest. Figure 4.7 shows the optimal amount of total wealth that should be allocated to housing. Note that PRE2015NOMWD has been left out as it perfectly overlaps with PRE2015, and POST2015 and POST2017 are almost identical to each other as well.

The calibrated risk aversion for housing (γ_{H}) falls between the risk aversions of singles and couples. The different curves for optimal housing are due to the marginal utility of housing in relation to the marginal utility of consumption and bequest. As wealth increases, the risk aversion for couples favours consumption more in relation to housing than that for single households. It is therefore optimal for single households to allocate a higher proportion than for couples. It might seem counter-intuitive, but this shoud be put into perspective: a couple household will have more assets available thus the average couple household will still allocate a higher dollar amount into housing compared with the average single household.

The literature is inconclusive as to whether retirees allocate more assets to housing in order to be eligible for an Age Pension or not. There is no evidence that the model considers means-test levels for optimal allocation to housing, such as whether the asset-test binds or not at t_0. This would be indicated by kinks in the housing allocation curve, where the remaining liquid wealth roughly equals the asset-test thresholds. The high proportion of housing allocation indicates that housing does play an important role in the Australian Age Pension. The different policy curves, however, shows that there is a possibility to plan housing allocation to maximise utility with respect to different policies, but not to Age Pension payments. After the changes to the income-test were introduced, the focus shifted even more to assets in the means-test and the optimal allocation then increased slightly.

4.6 Conclusion

In this chapter, the retirement model was calibrated against data from Australian Bureau of Statistics (2011) using the maximum likelihood method. Calibration was performed with respect to consumption and housing, and the estimated parameters

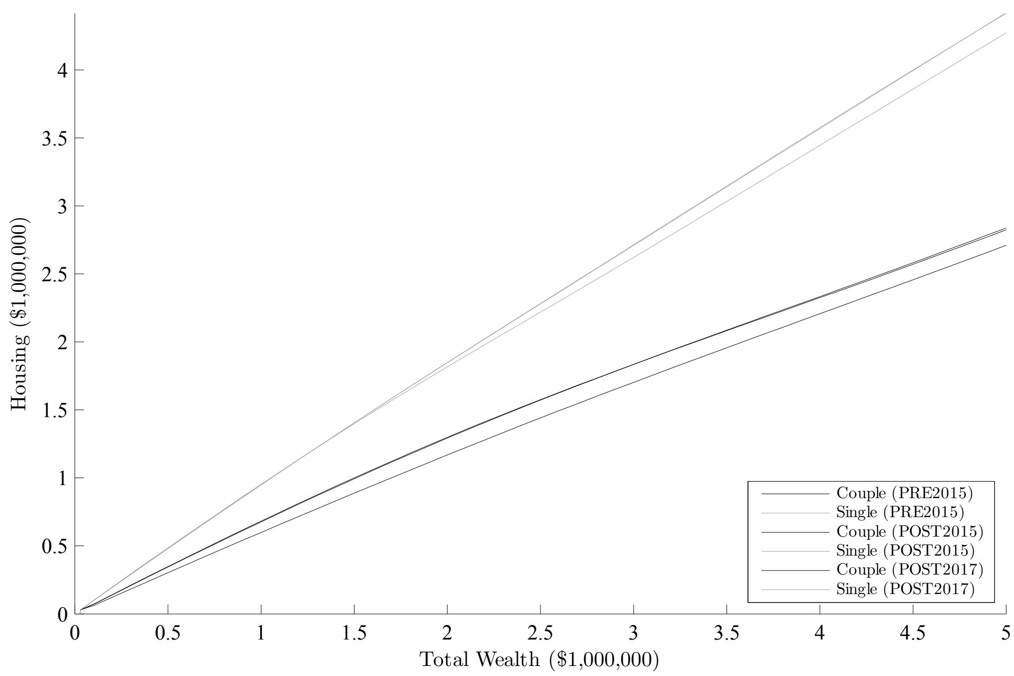

Figure 4.7: Optimal housing allocation given total wealth W for single and couple households under various Age Pension policies.

fall within proximity of related empirical research with the exception of bequest motives owing to differences in the model. A sensitivity test indicates that the model is rather robust, although individual changes in the risk aversion parameters lead to larger differences in optimal decisions.

The model was then solved for four different Age Pension policies and optimal decisions were analysed. The policies represent recent changes to the Age Pension and means-tests, where the first policy equals the one enforced during calibration. In general, the possibility of planning retirement decisions with respect to Age Pension is greater early in retirement but limited when minimum withdrawal rules exceed unconstrained optimal drawdown rates, especially for wealthier households. This tends to occur around ages 75-85, depending on wealth level. Only before this point, is it possible to plan withdrawals in order maximise utility, but these possibilities are almost nonexistent under the more recent policies. After this, optimal drawdown equals minimum withdrawal rates as it becomes a binding lower constraint for withdrawal.

Since drawdown of wealth is now replaced by deemed income, the assets are means-tested twice, which means risky asset allocation becomes more sensitive with respect to the asset-test. The means-test therefore plays a very important role when optimising risky asset allocation, as the Age Pension works as a 'buffer' for losses in financial assets. A potential loss can be (partly) offset by increased Age Pension payments, hence retirees who receive no or a partial Age Pension can accept more risk. This effect dies off as the minimum withdrawal rates bind, and the bequest motive becomes more important. In general, the changes in optimal risky asset allocation decreases with wealth, although it decreases over time for non-homeowners but increases for homeowners. The reason for the difference is that the home can be bequeathed, hence the part of liquid wealth in the bequest function is only a small part for homeowners, who then can accept more risk.

With regards to housing, it is optimal to allocate the majority of wealth to housing. This will allow the retiree to receive more partial Age Pension, while still being able to bequeath most of his/her wealth, and increases the expected utility in the long term. Couples have slightly lower preferences towards housing than single households, although this can be explained by the couple households having more accumulated wealth, thus the dollar amount invested in housing is still higher for an average couple household.

One surprising finding is that a retiree will receive more Age Pension over the course of their lifetime with the most recent policy rules. The high drawdown early in retirement would result in no Age Pension due to the income-test under the PRE2015 rules, while the new rules combined with the current historically low deeming rates will generate significant Age Pension payments from the same drawdown and wealth levels. This, in turn, affects both the decision for allocation in housing as well as risky investments. The government's goal of reducing incentives for maximising Age Pension payments and focusing on maximising total disposable income is, however, met - the current policy is not as sensitive to optimal withdrawal decisions in order to maximise Age Pension payments as the previous policies were.

Chapter 5

A Least-Squares Monte Carlo method for solving multi-dimensional expected utility models

5.1 Introduction

Stochastic control problems are at the heart of decision making under uncertainty and are critical in many areas such as finance, health, environment, and mining. In stochastic control problems there is always a choice to be made between model complexity, such as the number of state/control variables or stochastic factors, and computational cost. Analytical solutions are limited to problems with few stochastic factors with restrictions on the dynamics and dimensions, otherwise one has to revert to numerical methods. Partial differential equation methods suffer from the curse of dimensionality, and are practical up to two dimensions only. Numerical direct integration solutions, such as deterministic quadratures also suffer as the number of dimensions increases, but can sometimes handle more if we are willing to accept less precise solutions and longer computation times. Simulation methods are therefore favoured when the number of state variables and stochastic factors increases. The Stochastic Mesh method (Broadie and Glasserman, 2004) overcomes the dimensionality problem, but requires the transition densities of the stochastic factors to be known and suffers as the number of time steps increases. One simulation method that has received increasing interest among researchers is the Least-Squares Monte Carlo method (LSMC), due to its effectiveness in dealing with higher dimensions and because it imposes fewer restrictions on constraints and allows for flexibility in the dynamics of underlying stochastic processes. The idea is based on simulating random paths of the underlying stochastic variables over time and replacing the conditional expectation of the value function in the Bellman backward recursive solution of the stochastic control problem with an empirical least-squares regression estimate. The transition density of the underlying process is not even required to be known in closed form, which offers much more flexibility than alternative approaches.

LSMC was originally developed in Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001). The regression is generally performed on the state variables in order to approximate the value function. In the simpler case, where the state variable is exogenous (i.e. does not depend on control), the simulation and backward in time solution are rather straightforward. When considering endogenous state variables (i.e. affected by the control), however, the simulation becomes more complicated as the future states are affected by the unknown control. The extensions to LSMC that are of particular interest are methods where control variables are included in the regression basis functions[1] (Denault et al., 2013, 2017; Kharroubi et al., 2014). Kharroubi et al. (2014) allow for random control to be simulated and their algorithm (referred to as 'control randomisation') is the only theoretically justified LSMC algorithm with endogenous state variables. This algorithm provides additional benefits of parametric estimate in a feedback form of control, hence no solution grid for control is required (contrary to Denault et al. (2017)).

Naturally there are pitfalls with LSMC as well. Regression errors can accumulate over multiple time periods and can eventually blow up, and as the number of samples increases the algorithm becomes too computationally intensive. There are, however, methods to deal with such problems. The value function can either be based on the 'realised values' (Longstaff and Schwartz, 2001; Denault et al., 2013, 2017; Zhang et al., 2016) or on the 'regression surface' (Tsitsiklis and Van Roy, 2001; Denault et al., 2017). Although Denault et al. (2017) find little difference between 'realised

[1] A basis function is an element of a particular basis for a function space, where the full function space can be expressed as a linear combination of some chosen functions.

values' and the 'regression surface', it should be noted that the authors do not apply "true" realised values as the value function is interpolated with respect to the choice of controls. True realised values, which require re-simulation of paths after control is changed in order to calculate the value function, avoid regression errors to accumulate hence appear to be more stable over longer periods (Zhang et al., 2016) although with the trade-off of longer computation times. In addition, basis functions can often be difficult to find and can be highly problem specific. Incorrectly defined basis functions will quickly inflate the regression errors and blow up the solution of a multi-period stochastic control problem. This risk is especially high with regard to objective functions based on utility functions.

LSMC has been applied in many different fields, such as pricing American options (Longstaff and Schwartz, 2001; Tsitsiklis and Van Roy, 2001), mining and real options (Chen et al., 2015), electricity (Denault et al., 2013) and portfolio allocation (Brandt et al., 2005; Garlappi and Skoulakis, 2010; Zhang et al., 2016) to mention a few. Research with respect to certain issues in LSMC is very diverse, such as heteroskedasticity in the regression (Fabozzi et al., 2017), avoiding re-computing realised paths (Glasserman and Yu, 2004; Nadarajah et al., 2017; Nadarajah and Secomandi, 2017) or managing discontinuity in the basis function (Langrene et al., 2015).

With regard to problems involving utility functions however, there has been very limited research. Approximating a utility function with least-squares regression poses a number of challenges when the agent is risk averse due to very high second derivatives of the utility functions. Regression directly on the value function works only when samples of the state variable are restricted to a sub-domain rather than the full domain. If the full domain is used, then the fit of the regression will be unsatisfactory in parts of the domain due to the high curvature of the value function. This will not work in the case where a control can move the state variable over the full domain, such as a consumption problem where the decision can move the state of wealth from high to zero. It can work if the control has less influence over the change in the state variable, such as in portfolio allocation problems, as the allocation of assets will not result in as significant change in wealth compared to consumption. Even then authors acknowledge problems as volatility or risk aversion increases (Brandt et al., 2005; Denault et al., 2017).

Attempts have been made to resolve this issue, such as utilising Taylor series expansions around the value function (Brandt et al., 2005; Garlappi and Skoulakis, 2010) or by transformations. However, Taylor series expansions require the utility function to be differentiable, often with a minimum of four times and also add effort to compute the derivatives. Garlappi and Skoulakis (2010) apply an inverse utility function to the value function in order to perform the regression on the transformed value function, which does indeed improve the regression fit, but results in a 're-transformation bias' due to Jensen's inequality[2]. Such transformation also ignores any volatility of stochastic variables and underestimates risk. Denault et al. (2017) on the other hand only applies the transformation when interpolating in order to use a more coarse grid for state variables, but the use of a grid voids the purpose of 'true' LSMC. Zhang et al. (2016) suggest using basis functions with the independent variables transformed using a utility function, but if the domain covers a larger part of the curvature of the utility functions the regression errors will not be homogeneous, hence still result in a bias.

If the regression is carried out on the non-transformed value function, then control of disturbance terms is possible but the regression fit might be questionable. If the regression instead is carried out on the transformed value function, then the fit will most likely be better but optimal control related to disturbance will be unreliable. In this paper two methods are proposed to deal with the re-transformation bias, based on the characteristics of the problem, in order to account for difficulties in using LSMC with utility functions such as control of disturbance and re-transformation bias. In addition, a modification based on re-sampling state variables at each time step is suggested to the probabilistic numerical algorithm that combines dynamic programming with LSMC in Kharroubi et al. (2014, 2015) to improve the exploration of the state space, which further helps with the efficiency of the method in the case of expected utility models.

The chapter is structured as follows. In Section 5.2 the basic problem definition is stated, and Section 5.3 describes methods to avoid re-transformation bias. In Section 5.4 the LSMC algorithms are explained, while the accuracy of the algorithm

[2]Jensen's inequality states that for a random variable Z and a concave function ψ, $\psi(\mathbb{E}[Z]) \geq \mathbb{E}[\psi(Z)]$.

together with the methods to deal with the re-transformation bias are presented in Section 5.5. Finally, concluding remarks are in Section 5.6.

5.2 Problem definition

Let $t = 0, 1, ..., N$ correspond to equispaced points in time interval $[0, T]$ and $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{0\leq t\leq N}, \mathbb{P})$ be a filtered complete probability space where $\mathcal{F}_t$ represents the information available up to time t. We assume that all the processes introduced below are well defined and adapted to $\{\mathcal{F}_t\}_{t\geq 0}$. Let $\pi = (\pi_t)_{t=0,...,N}$ be a control taking value in an action space $\mathcal{A} \subseteq \mathbb{R}^d$, $Z = (Z_t)_{t=1,...,N} \in \mathcal{Z} \subseteq \mathbb{R}^d$ be a disturbance term with realisation z_t and $X^\pi = (X_t^\pi)_{t=0,...,N} \in \mathcal{X} \subseteq \mathbb{R}^d$ be a controlled state variable. We also assume that the evolution of the state variable is described by a transition function

$$X_{t+1}^\pi = \mathcal{T}_t\left(X_t^\pi, \pi_t, Z_{t+1}\right), \tag{5.1}$$

hence the state of the next period depends on the state of the current period, the control decision and the realisation of the disturbance term.

Now consider the standard discrete dynamic programming problem with the objective to maximise the expected value of total reward function

$$V_0(x) = \sup_\pi \mathbb{E}\left[\beta^N R_N(X_N^\pi) + \sum_{t=0}^{N-1} \beta^t R_t(X_t^\pi, \pi_t) \;\middle|\; X_0^\pi = x; \pi\right], \tag{5.2}$$

where R_N and R_t are reward functions satisfying the integrability conditions, and β is a time discount factor over a time step. This type of problem can be solved with backward recursion of the Bellman equation

$$\begin{aligned} V_t(x) &= \sup_{\pi_t}\left\{R_t(x, \pi_t) + \mathbb{E}\left[\beta V_{t+1}(X_{t+1}^\pi) \;\middle|\; X_t^\pi = x; \pi_t\right]\right\}, t = N-1, ..., 0, \\ V_N(x) &= R_N(x). \end{aligned} \tag{5.3}$$

The solution of such a problem is often not possible to find analytically and numerical methods are required. As the number of state variables, stochastic processes, or control variables increases, the numerical solution quickly becomes very computationally expensive. In the standard case, where state is not affected by control, the idea behind utilitising the LSMC method is to approximate the conditional expectation in Equation (5.3)

$$\Phi_t(X_t) = \mathbb{E}\left[\beta V_{t+1}(X_{t+1})|X_t\right], \tag{5.4}$$

by a regression scheme with independent variables X_t, and response variable $\beta V_{t+1}(X_{t+1})$. The approximation of the function is then denoted as $\widehat{\Phi}_t$. However, if the state is affected by control, then control randomization is required and the conditional expectation

$$\Phi_t(X_t^\pi, \pi_t) = \mathbb{E}\left[\beta V_{t+1}(X_{t+1}^\pi)|X_t^\pi; \pi_t\right], \tag{5.5}$$

is estimated by regression on X_t^π and randomised π_t (Kharroubi et al., 2014).

For ease of notation the superscript π on the state variable is now dropped.

5.3 Transformation of utility

One of the difficulties with LSMC is to select correct basis functions for regression estimate of conditional expectation $\Phi_t(X_t, \pi_t)$. Commonly used basis functions include polynomials such as Chebyshev, Hermite, Laguerre and Legendre (Table 5.1) or one can use Appell polynomials as the generalized form of many standard polynomials (Novikov and Shiryaev, 2005). Unless the function that is being approximated is convex/concave or smooth, such as piecewise linear, the basis functions might only work locally. Increasingly complex (e.g., higher order) basis functions can be used with increasing number of sample paths to improve the accuracy, but this comes with a computational cost. However, when reward functions $R_t(x, \cdot)$ in (5.3) are based on the standard utility functions, such as Constant Relative Risk Aversion (CRRA) $U(x) = x^\gamma/\gamma$ or Hyperbolic Absolute Risk Aversion (HARA) $U(x) = (x-a)^\gamma/\gamma$, the basis functions do not produce accurate solution unless constrained locally. A model based on utility functions will have a value function with a similar shape, hence the same problems will arise when fitting the basis functions to $V_{t+1}(X_{t+1})$ as if fitting it to a utility function. A more effective solution is to perform a transformation based on the utility function, and account for the re-transformation bias.

Table 5.1: Definition of common polynomials used as basis functions up to the nth order.

Polynomial	$f_0(x)$	$f_1(x)$	$f_n(x)$
Chebyshev	1	x	$2xf_{n-1}(x) - f_{n-2}(x)$
Hermite	1	x	$(-1)^n e^{-x^2} \frac{d^n}{dx^n}(e^{-x^2})$
Laguerre	1	$1-x$	$\frac{(2(n-1)+1-x)f_{n-1}(x)-(n-1)f_{n-2}(x)}{n}$
Legendre	1	x	$\frac{1}{2^n n!}\frac{d^n}{dx^n}(x^2-1)n$

Define a utility function $U : \mathbb{R} \to \mathbb{R}$ to be an increasing, monotonic and concave function. Consider a stochastic control problem (5.3) where the reward is based on this utility function. Then a value function $V_t(x)$ equals the utility function at time T and will have a similar shape for $t < N$. Such a function is difficult to fit with linear regression due to the extreme curvature of common utility functions. For example, consider a CRRA utility function $U(x) = x^\gamma/\gamma$ with a risk aversion parameter $\gamma < 0$. The first problem arises when fitting the regression to low values of x, since when $x \to 0$ then $U(x) \to -\infty$ and therefore no intercept exists. This could be avoided by using fractional polynomials in the basis functions, which are polynomials with fractional exponents. These polynomials (or independent variables with negative exponents) tend to approximate the utility function shape better. Unfortunately, fitting such a model is non-trivial when the utility function is not CRRA or the utility function exerts any kind of piecewise behaviour in relation to the state variables. If a transformation is applied to either decrease the non-linearity in the utility function or to deal with non-normality and heteroskedastic residuals, then Jensen's inequality results in an incorrect projection when the regression is transformed back since $\mathbb{E}[U(\cdot)] \leq U(\mathbb{E}[\cdot])$ due to U being concave. The solution will therefore be biased. In addition to this, any control variables that relate to disturbance terms (such as allocation between risky and risk-free assets in wealth allocation problems) will be biased because risk is underestimated.

To improve the approximation of the value function with the least-squares regression we propose that the regression is performed on the *transformed* value function and adjusted for the inverse transformation bias with a bias correction function to account for Jensens's inequality. The value function transformed using the inverse of the utility function will have less non-linearity and will allow for an intercept, hence it will have better fit with linear regression (although non-linear independent variables might still be required). Specifically, we proceed as follows.

Define a transformation $H^{-1} : \mathbb{R} \to \mathbb{R}$ and the inverse ('re-transformation') $H : \mathbb{R} \to \mathbb{R}$ such that $H^{-1}(H(x)) = x$. It is implied that state variables still depend on control. Let $\mathbf{L}(X_t, \pi_t)$ be a vector of basis functions and $\mathbf{\Lambda}_t$ the corresponding regression coefficients vector, such that

$$\mathbb{E}\left[H^{-1}(\beta V_{t+1}(X_{t+1}))|X_t; \pi_t\right] = \mathbf{\Lambda}_t'\mathbf{L}(X_t, \pi_t). \tag{5.6}$$

If M independent Markovian paths of state and control variables are simulated, one can consider the ordinary linear regression

$$\begin{aligned} &H^{-1}(\beta V_{t+1}(X_{t+1}^m)) = \mathbf{\Lambda}_t'\mathbf{L}(X_t^m, \pi_t) + \epsilon_t^m, \\ &\epsilon_t^m \overset{iid}{\sim} F_t(\cdot), \quad \mathbb{E}[\epsilon_t^m] = 0, \quad \text{var}[\epsilon_t^m] = \sigma_t^2, \quad m = 1, ..., M \end{aligned} \tag{5.7}$$

to estimate the regression coefficients as

$$\widehat{\mathbf{\Lambda}}_t = \arg\min_{\mathbf{\Lambda}_t} \sum_{m=1}^{M} \left[H^{-1}(\beta V_{t+1}(X_{t+1}^m)) - \mathbf{\Lambda}_t'\mathbf{L}(X_t^m, \pi_t^m)\right]^2. \tag{5.8}$$

It is well known that the estimator $\widehat{\mathbf{\Lambda}}_t$ is the best linear unbiased estimator which is also consistent and asymptotically normally distributed. If the disturbances ϵ_t^m are normally distributed, then this estimator is the maximum likelihood estimator and asymptotically efficient. Moreover, if disturbances ϵ_t^m are heteroscedastic (have different variance), this estimator remains unbiased, consistent, and asymptotically normally distributed but no longer efficient; see for example (Greene, 2008, chapter 8).

Our objective is to estimate $\Phi_t(X_t^\pi, \pi_t) = \mathbb{E}\left[\beta V_{t+1}(X_{t+1}^\pi)|X_t^\pi; \pi_t\right]$ that can be expressed as

$$\Phi_t(X_t, \pi_t) := H^B(\mathbf{\Lambda}_t'\mathbf{L}(X_t, \pi_t)) = \int H(\mathbf{\Lambda}_t'\mathbf{L}(X_t, \pi_t) + \epsilon_t)dF_t(\epsilon_t), \tag{5.9}$$

where $F_t(\epsilon_t)$ is the distribution of disturbance term ϵ_t. Obviously, in general, naive estimation

$$\widehat{H}^B(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t,\pi_t)) = H(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t,\pi_t)) \tag{5.10}$$

will be neither unbiased nor consistent (even if we know the true parameters $\mathbf{\Lambda}_t$) unless the transformation is linear.

If a specific distribution is assumed for ϵ_t, then the integration in (5.9) can be performed (in closed form for some cases). Otherwise, the empirical distribution of residuals

$$\widehat{\epsilon}_t^m = H^{-1}(\beta V_{t+1}(X_{t+1}^m)) - \widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t^m,\pi_t^m), \tag{5.11}$$

can be used to perform the required integration as proposed in (Duan, 1983) leading to the following estimate.

Smearing Estimate:

$$\begin{aligned}\widehat{H}^B(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t,\pi_t)) &= \int H(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t,\pi_t)+\epsilon_t)d\widehat{F}_M(\epsilon_t)\\ &= \frac{1}{M}\sum_{m=1}^{M} H(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t,\pi_t)+\widehat{\epsilon}_t^m),\end{aligned} \tag{5.12}$$

where $\widehat{F}_M(\epsilon_t)$ is the empirical distribution function of the estimated residuals (see Appendix B for details).

If heteroskedasticity is present in the regression with respect to state and control variables, a method that accounts for the heteroskedasticity is required. In this case the conditional variance can be modelled as a function of covariates,

$$\mathrm{var}[\epsilon_t|X_t,\pi_t] = [\Omega(\boldsymbol{\mathcal{L}}_t'\mathbf{C}(X_t,\pi_t))]^2, \tag{5.13}$$

where $\Omega(\cdot)$ is some positive function, $\boldsymbol{\mathcal{L}}_t$ is the vector of coefficients and $\mathbf{C}(X_t,\pi_t)$ is a vector of basis functions. There are various standard ways to find estimates $\widehat{\boldsymbol{\mathcal{L}}}_t$, the one we use in this chapter is based on squared residuals of the ordinary least squares method as outlined in Appendix C. Then, one can use the Smearing Estimate with Controlled Heteroskedasticity proposed in (Zhou et al., 2008) and defined as follows.

Smearing Estimate with Controlled Heteroskedasticity:

$$\widehat{H}^B(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t,\pi_t)) = \frac{1}{M}\sum_{m=1}^{M} H\left(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t,\pi_t) + \Omega(\widehat{\boldsymbol{\mathcal{L}}}_t'\mathbf{C}(X_t,\pi_t))\frac{\widehat{\epsilon}_t^m}{\Omega(\widehat{\boldsymbol{\mathcal{L}}}_t'\mathbf{C}(X_t^m,\pi_t^m))}\right). \tag{5.14}$$

Here, it is also common to replace $\widehat{\mathbf{\Lambda}}_t$ with the weighted least squares estimator that can be found after estimation of $\Omega(\cdot)$, see for example two-step procedure in (Greene, 2008, chapter 8).

It should be noted that an alternative would be to use Generalised Linear Models, where no transformation of the value function is required and which allow for heteroskedasticity through a link function. However, Generalised Linear Models are reported to be quite imprecise when the error distribution assumptions are inaccurate (Baser, 2007), or if the distribution family is misspecified. For a more flexible approach with fewer restrictions we prefer to use Smearing Estimate if no control of disturbance terms is required, and Smearing Estimate with Controlled Heteroskedasticity if control variables are related to disturbance terms in the model. These methods also have the additional advantage that the utility function does not need to be differentiable or continuous, as long as a transformation that roughly represents the shape of the value function can be found.

5.4 LSMC algorithm

In this section we describe the LSMC algorithms for the exogenous state and the endogenous state with control randomisation. In addition, we use an example with Bermudan options to benchmark the re-transformation method with bias correction against the standard LSMC method.

5.4.1 Basic algorithm with exogenous state

The basic exogenous state LSMC commonly used in research literature, presented below in Algorithm 2, is based on two parts: a forward simulation and a backward solution with optimal control. First, random state paths $X_t^m, m = 1, ..., M$, are generated which are affected by random disturbances, following the state evolution in Equation (5.1) for $t = 0, ..., N$. The problem is then solved in a dynamic programming fashion, where the reward function is first evaluated at time $t = N$. Then, starting at $t = N - 1$ we find the optimal control for each t by regressing the value function at $t + 1$ on the state variables at t. Once a decision has been made, the value function is updated with the outcome and calculations are repeated for $t - 1$ until we find value function at 0.

Algorithm 2 LSMC for exogenous state

[Forward simulation]
for $t = 0$ **to** N **do**
    for $m = 1$ **to** M **do**
        [Simulate random path]
        if $t = 0$ **then**
            $X_t^m := S_0$
        else
            $X_t^m := \mathcal{T}_t(X_{t-1}^m, z_t)$
        end if
    end for
end for

[Backward solution]
for $t = N$ **to** 0 **do**
    if $t = N$ **then**
        $\widehat{V}_t(\boldsymbol{X}_t) := R_N(\boldsymbol{X}_t)$
    else if $t < N$ **then**
        [Regress transformed value function on state variables]
        $\widehat{\mathbf{\Lambda}}_t := \arg\min_{\mathbf{\Lambda}_t} \sum_{m=1}^{M} \left[\mathbf{\Lambda}_t' \mathbf{L}(X_t^m) - H^{-1}(\beta \widehat{V}_{t+1}(X_{t+1}^m))\right]^2$
        Find bias corrected transformation $H^B(\widehat{\mathbf{\Lambda}}_t' \mathbf{L}(X_t))$
        [Approximate conditional expectation]
        $\widehat{\Phi}_t(X_t) := H^B(\widehat{\mathbf{\Lambda}}_t' \mathbf{L}(X_t))$
        for $m = 1$ **to** M **do**
            [Optimal control]
            $\pi_t^*(X_t^m) := \arg\sup_{\pi_t \in A} \left\{R_t(X_t^m, \pi_t) + \widehat{\Phi}_t(X_t^m)\right\}$
            $\widehat{V}_t(X_t^m) := R_t(X_t^m, \pi_t^*(X_t^m)) + \beta \widehat{V}_{t+1}(X_{t+1}^m)$
        end for
    end if
end for

To illustrate the LSMC algorithm, we start with a basic exogenous state version of LSMC applied to pricing standard Bermudan options in the same manner as in Longstaff and Schwartz (2001). These options can only be exercised at pre-specified dates and at maturity. Although an option that can be exercised prior to maturity does not have a truly exogenous state (as the state changes if the option is exercised), it can still be written in such a way that it can be presented in the current framework and the exogenous state algorithm can be used. Since Algorithm 2 shows the general algorithm for the exogenous state, some minor changes are required in order to use it for Bermudan options and will be discussed further down.

Since utility functions are not used in the pricing, no transformation of the value function is required (hence the value function estimate will not suffer from a re-transformation bias other than possible bias from regression errors). Consider a put option as in the original paper (Longstaff and Schwartz, 2001). Let $t = 0, 1, ..., N$ correspond to the pre-specified equidistant exercise dates obtained by dividing the option maturity T into N steps of length $\delta t = T/N$. The option underlying asset price S_t evolves (under the so-called risk-neutral process appropriate for valuation of option fair value) as

$$S_{t+1} = S_t e^{(r-\sigma^2/2)\delta t + \sigma Z_t \sqrt{\delta t}}, \quad Z_t \overset{i.i.d}{\sim} \mathcal{N}(0,1), \tag{5.15}$$

where r is risk free interest rate, σ is volatility and $\mathcal{N}(0,1)$ is the standard normal distribution. The discounting factor for each time period is $\beta = e^{-r\delta t}$. The control variable takes on two values, $\pi_t \in \{0, 1\}$, which represent continuation if $\pi_t = 0$ or immediate exercising if $\pi_t = 1$. The state variable consists of the current asset price and an absorbing state, $X_t = \{S_t, \Delta\}$, where Δ indicates that the option has already been exercised. The transition probabilities for the absorbing state are

$$\Pr[X_{t+1} = \Delta | X_t = \Delta] = \Pr[X_{t+1} = \Delta | X_t = S_t, \pi_t = 1] = 1, \tag{5.16}$$

while the transition probability $\Pr[X_{t+1} \in dS_{t+1} | X_t = S_t, \pi_t = 0]$ corresponds to the process for the asset price. Any other transitions cannot occur, such as moving to the absorbing state if the option is not exercised, hence the remaining transition probabilities are zero.

The terminal reward depends on the moneyness of the option at expiration assuming it has not been exercised,

$$R_N(X_N) = \begin{cases} \max(0, K - X_N), & \text{if } X_N \neq \Delta, \\ 0, & \text{if } X_N = \Delta, \end{cases} \tag{5.17}$$

where K is the option strike price. The reward function at time t

$$R_t(X_t, \pi_t) = \begin{cases} \max(0, K - X_t), & \text{if } X_t \neq \Delta \text{ and } \pi_t = 1, \\ 0, & \text{otherwise}, \end{cases} \tag{5.18}$$

only provides reward (a payoff) if the option has not been exercised earlier and the decision is to exercise the option immediately. The solution of the problem starts by evaluating the payoff at time N. The cash flow at this point is stored as $V_N(X_N)$ and at each previous time the decision to exercise the option or wait is determined by comparing the immediate payoff $\max(0, K - X_t)$ with the continuation value $\mathbb{E}[e^{-r\delta t}V_{t+1}(X_{t+1})|X_t]$. The continuation value is the discounted expected value if the option is not exercised at time t. The estimation of the continuation value is done by regressing the realised cash flows of $V_{t+1}(X_{t+1})$ at time $t+1$ discounted to t (if not exercised) on a vector $\mathbf{L}(X_t)$ of basis functions of the state variable

$$\mathbb{E}[e^{-r\delta t}V_{t+1}(X_{t+1})|X_t] = \mathbf{\Lambda}_t'\mathbf{L}(X_t), \tag{5.19}$$

in order to approximate the conditional expectation $\Phi_t(X_t) = \mathbb{E}\left[\beta V_{t+1}(X_{t+1})|X_t\right]$. Note that the state variable, if not exercised, is exogenous and not affected by any control, hence the transition function in Equation (5.1) is simplified and depends only on the previous state and the outcome of the disturbance term(s). The optimal control can then be written as

$$\pi_t^*(X_t) := \arg\max_{\pi_t} \left\{R_t(X_t, \pi_t) + (1 - \pi_t)\Phi_t(X_t)\right\}, \tag{5.20}$$

since if the option is now exercised the continuation value will not be received, which is reflected by $(1 - \pi_t)$ in front of the conditional expectation approximation, and is the same as all future rewards are zero as the state will transition to the absorbing state Δ. Equation (5.20) therefore replaces line 18 in Algorithm 2. The effect of the decision is recorded in the realised value of

$$\widehat{V}_t(X_t) = R_t(X_t, \pi_t^*(X_t)) + (1 - \pi_t^*(X_t))e^{-r\delta t}\widehat{V}_{t+1}(X_{t+1}), \tag{5.21}$$

which replaces line 19. If the option is exercised, the realised value equals the reward for the current period, and if it is not exercised it equals the present value of future rewards. The full objective function for the Bermudan option problem, which originally is an optimal stopping problem, then leads to the same optimal stochastic control problem as Equation (5.2).

As a numerical example, consider the Bermudan option when $S_0 = 36$, $K = 40$, $r = 0.06$, $\sigma = 0.2$, $T = 1$ and $N = 12$ (for results, see Table 5.2). The basis functions are based on ordinary polynomials up to the fourth order of the state variable, and the price is benchmarked against the Binomial Tree method. First, the problem is solved using the standard LSMC (column $V^{(0)}$). By using a log transformation of the value function, i.e. $H(x) = e^x$, the heteroskedasticity in the residual errors is reduced and we get a better regression estimation, although the bias still remains when re-transformed (column $V^{(1)}$ where naive estimation (5.10) is used). Finally, we correct for the bias after re-transformation using Smearing Estimate (5.12). This log transformation with the bias correction results in a more accurate price even with fewer sample paths (column $V^{(2)}$).

Table 5.2: Price and standard error of Bermudan option estimated using standard LSMC ($V^{(0)}$), LSMC with log transformation of the value function without bias correction ($V^{(1)}$) and LSMC with log transformation of value function and bias correction ($V^{(2)}$) using Smearing Estimate. The results are based on M sample paths, 20 independent repetitions (iterations), and the basis functions are ordinary polynomials up to the 4th order. The 'exact' price obtained by the finite difference method is \$4.3862.

M	$V^{(0)}$	$V^{(1)}$	$V^{(2)}$
1,000	4.4984 (0.032)	4.4336 (0.038)	4.4054 (0.039)
10,000	4.4616 (0.007)	4.4161 (0.007)	4.3962 (0.008)
100,000	4.4457 (0.003)	4.4048 (0.004)	4.3857 (0.004)

5.4.2 Endogenous state and random control

Algorithm 2 presented in previous section is the very basic case where optimal decisions do not really affect the evolution of the state variable, with the exception of reaching the absorbing (exercised) state. To extend it to the case with an endogenous state, we adopt the discretised version of the control randomisation technique and LSMC algorithm with realised values from Kharroubi et al. (2014), which is the only theoretically justified LSMC algorithm with endogenous state variables. The algorithm is also based on forward simulation and backward solution with optimal control, with the main difference that random state paths in the forward simulation are affected by both a random control and random disturbance. The regression to estimate the conditional expectation then includes the random control in the basis functions. Kharroubi et al. (2014) present two alternative versions of the control randomisation algorithm: the one that uses the regression surface to update the value function,

$$\widehat{V}_t(X_t) = R_t(X_t, \pi_t^*(X_t)) + \widehat{\Phi}_t(X_t, \pi_t^*(X_t)), \tag{5.22}$$

and another one that uses the realised value function,

$$\widehat{V}_t(X_t) = R_t(X_t, \pi_t^*(X_t)) + \beta \widehat{V}_{t+1}(X_{t+1}). \tag{5.23}$$

The difference is that the first method is a value function iteration (VFI), while the second is a policy function iteration (PFI). The PFI requires a recalculation of the sample paths for $t+1$ to T after each iteration backwards, as the optimal control affects the future state variables hence changes the simulated paths. This effect is already estimated in the VFI version with $\Phi_t(X_t, \pi_t)$, hence no recalculation is necessary. Longstaff and Schwartz (2001) argue in favour of realised value function (5.23), while Tsitsiklis and Van Roy (2001) use the regression surface (5.22). Denault et al. (2017) notices no difference between the two, but only use a very local regression which does not include the endogenous state variable. Due to the recalculation requirement with PFI the computational complexity grows quadratically compared with linear growth for VFI. However, PFI tends to accumulate much less regression errors over time and from experience this method is much more suitable for problems prone to regression errors when the number of time periods increases.

Forward simulation

The forward simulation in the case of endogenous state with control randomisation is more delicate than in the case of exogenous state variables, and deserves a special discussion below. The objective of the forward simulation is to generate enough information (sample paths) such that the conditional expectation of state and control variables can be estimated.

The algorithm is intended to be used with a known starting state x_0 where each simulated path X_t^m, $t = 0, ..., N$, is subject to the random control $\widetilde{\pi}_t^m$ and disturbances in the diffusion process, given by the evolution in Equation (5.1). However, if the optimal control π_t^* tends to take on either a high or low value in the control domain $\mathcal{A}$ this will cause a problem. On the one hand, if random control is uniformly distributed then the simulated paths will end up in a sub-domain very different from the one if optimal control is used at each time step. On the other hand, if the typical range of the optimal control is known, a distribution to reflect this could be used to better simulate the randomised control where more paths would end up in the same sub-domain as if the optimal control was applied. While the former would lead to difficulties in the regression estimation due to lack of state sample paths in sections of the state domain, the latter would do the same due to lack of control samples outside the expected range (hence lack samples in the full control domain $\mathcal{A}$).

A better approach would be to simulate state and control for the full domain to ensure a better fit for the regression. This can be achieved by using a random state each time t that is independent of decisions and disturbance for $0, ..., t-1$, as the simulated paths are recalculated after each time period anyway. Denote $\widetilde{X}_t$ as the state variable implied by transition function $\mathcal{T}_{t-1}(X_{t-1}, \pi_{t-1}, z_t)$, where X_t is an independent random sample from the state variable domain at each time step. If the state variable X_t would be simulated using the transition function as a path for the full period $t = 0, ..., T$, then $\widetilde{X}_{t+1} = X_{t+1}$ would hold. The logical steps of this procedure are summarized by Algorithm 3 below. Each X_t is simulated independently of the previous state, which allows the algorithm to spread samples over the full domain each time period to avoid the pitfalls described. This will explore the space better and the reason for this will become apparent in Algorithm 4. In the algorithm, *Rand* corresponds to random sampling from some distribution that could be designed for the specific problem. These can be, e.g., uniform distributions for X_t and π_t, while distribution for z_t is model specific.

Algorithm 3 Forward simulation

for $t = 0$ **to** $N-1$ **do**
  for $m = 1$ **to** M **do**
    [Simulate random samples]
    $X_t^m := Rand \in \mathcal{X}$ ▷ State
    $\widetilde{\pi}_t^m := Rand \in \mathcal{A}$ ▷ Control
    $z_{t+1}^m := Rand \in \mathcal{Z}$ ▷ Disturbance
    [Compute the state variable after control]
    $\widetilde{X}_{t+1}^m := \mathcal{T}_t(X_t^m, \widetilde{\pi}_t^m, z_{t+1}^m)$ ▷ Evolution of state
  end for
end for

Backward optimisation

After the forward simulation step is completed, the problem is now solved with the backward induction (Algorithm 4), similarly to Algorithm 2. The conditional expectation $\Phi_t(\cdot)$ of the value function at time $t+1$ is estimated with a regression function, and the optimal decision is found by maximising the sum of the reward function and the approximated value function. Once the optimal decision has been found, the sample paths $t+1, ..., N$ are recalculated with the new optimal control and the corresponding value function for the realisation of the paths is stored to be used in the next iteration.

Algorithm 4 Backward solution (Realised value)

1: **for** $t = N$ **to** 0 **do**
2:   **if** $t = N$ **then**
3:     $\widehat{V}_t(\widetilde{\boldsymbol{X}}_t) := R_N(\widetilde{\boldsymbol{X}}_t)$
4:   **else if** $t < N$ **then**
    [Regression of transformed value function]
5:     $\widehat{\boldsymbol{\Lambda}}_t := \arg\min_{\boldsymbol{\Lambda}_t} \sum_{m=1}^{M} \left[\boldsymbol{\Lambda}_t' \mathbf{L}(X_t^m, \widetilde{\pi}_t) - H^{-1}(\beta \widehat{V}_{t+1}(\widetilde{X}_{t+1}^m))\right]^2$
6:     Find bias corrected transformation $H^B(\widehat{\boldsymbol{\Lambda}}_t' \mathbf{L}(X_t, \widetilde{\pi}_t))$
    [Approximate conditional expectation]
7:     $\widehat{\Phi}_t(X_t, \widetilde{\pi}_t) := H^B(\widehat{\boldsymbol{\Lambda}}_t' \mathbf{L}(X_t, \widetilde{\pi}_t))$
8:     **for** $m = 1$ **to** M **do**
9:       $\widehat{X}_t^m := \widetilde{X}_t^m$
      [Optimal control]
10:       $\pi_t^*(\widehat{X}_t^m) := \arg\sup_{\pi_t \in A} \left\{R_t(\widehat{X}_t^m, \pi_t) + \widehat{\Phi}_t(\widehat{X}_t^m, \pi_t)\right\}$
      [Update value function with optimal paths]
11:       $\widehat{V}_t(\widehat{X}_t^m) := R_t(\widehat{X}_t^m, \pi_t^*(\widehat{X}_t^m))$
12:       $\widehat{X}_{t+1}^m := \mathcal{T}_t(\widehat{X}_t^m, \pi_t^*(\widehat{X}_t^m), z_t^m)$
13:       **for** $t_j = t+1$ **to** $N-1$ **do**
14:         $\widehat{V}_t(\widehat{X}_t^m) := \widehat{V}_t(\widehat{X}_t^m) + \beta^{t_j - t} R_{t_j}(\widehat{X}_{t_j}^m, \pi_{t_j}^*(\widehat{X}_{t_j}^m))$
15:         $\widehat{X}_{t_j+1}^m := \mathcal{T}_t(\widehat{X}_{t_j}^m, \pi_{t_j}^*(\widehat{X}_{t_j}^m), z_{t_j}^m)$
16:       **end for**
17:       $\widehat{V}_t(\widehat{X}_t^m) := \widehat{V}_t(\widehat{X}_t^m) + \beta^{N-t} R_N(\widehat{X}_N^m)$
18:     **end for**
19:   **end if**
20: **end for**

Note that at terminal time $t = N$, where no decision is allowed, $\widetilde{X}_N = X_N$ holds true, which is why $\widetilde{X}_N$ is used on line 3 in Algorithm 4 (hence no need to simulate X_N which is reflected in Algorithm 3 as it stops at $N-1$). Furthermore, on line 9 and 10 the state after control $\widetilde{X}_t$ is used, in order to estimate the value function for the current period t on line 14. This way π_t^* has been found and the value function for $\widetilde{X}_t$ rather than X_t has already been prepared so it can be used directly for the next iteration.

The loop starting at line 13 is crucial for multi-period stochastic control problems with utility functions (unless the basis functions are correct for all periods). It updates the forward simulation at each backward step with the optimal decision, similar to Longstaff and Schwartz (2001). It therefore uses the realised value function rather than the regression surface (VFI and PFI methods in Equation (5.22) and (5.23)) and takes advantage of the tower property of conditional expectations[3]. This step helps significantly with the accuracy of the approximation as the time horizon extends, and avoids (or at least reduces) the risk of the solution blowing up by limiting the accumulation of regression errors. Compare this with Algorithm 5, which is the equivalent of Algorithm 4 but based on the regression surface (VFI approach in Equation (5.22)) rather than realised values (PFI approach in Equation (5.23)). The algorithm does not update the forward simulation at each pass, hence is faster but might pile regression errors in the value function. However, in Algorithm 4, the function for the optimal control $\pi_t^*(\widehat{X}_t)$ is solved each time period during the forward loop on line 13 that makes the algorithm more computationally expensive. By storing the optimal control for each state sample, rather than just the regression coefficients, the optimal control can instead be interpolated for each t. This significantly speeds up the solution, especially as the number of dimensions/controls increases.

Algorithm 5 Backward solution (Regression surface)

for $t = T$ **to** 0 **do**
  if $t = T$ **then**
    $\widehat{V}_t(\widetilde{\mathbf{X}}_t) := R_N(\widetilde{\mathbf{X}}_t)$
  else if $t < T$ **then**
    [Regression of transformed value function]
    $\widehat{\mathbf{\Lambda}}_t := \arg\min_{\mathbf{\Lambda}_t} \sum_{m=1}^{M} \left[\mathbf{\Lambda}_t'\mathbf{L}(X_t^m, \widetilde{\pi}_t) - H^{-1}(\beta \widehat{V}_{t+1}(\widetilde{X}_{t+1}^m))\right]^2$
    Find bias corrected transformation $H^B(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t, \widetilde{\pi}_t))$
    [Approximate conditional expectation]
    $\widehat{\Phi}_t(X_t, \widetilde{\pi}_t) := H^B(\widehat{\mathbf{\Lambda}}_t'\mathbf{L}(X_t, \widetilde{\pi}_t))$
    for $m = 1$ **to** M **do**
      $\widehat{X}_t^m := \widetilde{X}_t^m$
      [Optimal control]
      $\pi_t^*(\widehat{X}_t^m) := \arg\sup_{\pi_t \in A} \left\{R_t(\widehat{X}_t^m, \pi_t) + \widehat{\Phi}_t(\widehat{X}_t^m, \pi_t)\right\}$
      $\widehat{V}_t(\widehat{X}_t^m) := R_t(\widehat{X}_t^m, \pi_t^*(\widehat{X}_t^m)) + \widehat{\Phi}_t(\widehat{X}_t^m, \pi_t^*(\widehat{X}_t^m))$
    end for
  end if
end for

[3]The tower property states that when conditioning twice, with respect to nested σ-algebras, the smaller amount of information always prevails such that $\mathbb{E}[\mathbb{E}[Z|\mathcal{F}_{t+1}]|\mathcal{F}_t] = \mathbb{E}[Z|\mathcal{F}_t]$

If a fixed starting point X_0 is desired as in the original algorithm, rather than a range of potential starting points, then line 9 would be replaced at $t = 0$ with $\widehat{X}_0^m := X_0$.

5.4.3 Upper and lower bounds

The value function from Algorithm 4 is already a lower bound (up to the Monte Carlo error), as replacing the supremum in Equation (5.3) with the estimated optimal control yields a lower bound by definition of the supremum (Aïd et al., 2014). Similarly, an approximate upper bound can be found by using the optimal control estimate π_t^* from Algorithm 4, but on line 11 in Algorithm 5 where the realised value function is replaced by the regression surface. Given that estimator $\hat{\Phi}_t(X_t, \pi_t)$ of the approximated conditional expectation in Equation (5.5) is unbiased, this results in an upper bound up to the Monte Carlo and regression error (Aïd et al., 2014).

Even if the transformation method minimises the regression error, the error will always be present. Given a concave utility function these errors will be biased downwards. Since this upper bound is based on the regression surface, rather than realised values, it will accumulate significant regression bias over time and often leads to lower value than the lower bound. An alternative approach is to use an upper bound based on the expected change of the stochastic variables, as this will always be equal or larger according to Jensen's inequality given a concave utility function. As the lower and upper bound now will suffer equally from the regression bias it represents a more realistic upper bound of the solution. This only holds true for a concave utility function – if the function is convex the inequality changes direction.

5.5 Accuracy of solution

In this section we examine the impacts of Smearing Estimate (5.12), Smearing Estimate with Controlled Heteroskedasticity (5.14) and naive estimate without transformation bias correction (5.10) on the accuracy of the LSMC numerical solution. Two basic models based on CRRA utility are considered: optimal consumption, and optimal consumption and risky asset allocation. The models have closed form solutions presented in Appendix D which are used for benchmarking of our numerical methods. The chosen model parameter values correspond to the ones causing problems with numerical solutions in Denault et al. (2017) in the case of high risk aversion and volatility.

In both examples, we use 10 000 sample paths where state and control were sampled from uniform distributions, and the disturbance term from a normal distribution. The basis functions are ordinary polynomials up to the 4th order in both transformed state and control variables, including mixed terms. The transformation used is based on the exponent and a CRRA utility function, such that

$$H^{-1}(x) = \ln\left[(\gamma x)^{1/\gamma}\right]. \tag{5.24}$$

Note that the examples below do not include the standard LSMC case, and use re-sampling in the forward simulation (Algorithm 3). The reason is simply because the standard LSMC method is not stable and the solution either blows up, or optimal control equals full or zero consumption due to bad regression fit if re-sampling in the forward simulation step is not used.

5.5.1 Consumption model

Consider a typical simple model where the agent receives utility by consuming a proportion $\pi_t := \alpha_t \in [0,1]$ of the endogenous state variable wealth X_t each period $t = 0, ..., 9$, hence terminal time $T = N = 9$ resulting in 10 evaluations. The utility at time t is $R_t(X_t, \pi_t) = (\alpha_t X_t)^\gamma/\gamma$ and utility of wealth at terminal time is $R_N(X_N) = (\alpha_N X_N)^\gamma/\gamma$ with risk aversion $\gamma = -10$. Wealth change between periods is based on a stochastic return $Z \sim \mathcal{N}(\mu, \sigma^2)$ with drift $\mu = 0.1$ and standard deviation $\sigma = 0.2$, such that the transition to the wealth at $t+1$ is

$$X_{t+1} = \mathcal{T}_t(X_t, \pi_t, Z_{t+1}) := X_t(1-\alpha_t)e^{Z_{t+1}}. \tag{5.25}$$

The closed-form solution for optimal consumption is then

$$\alpha_t = \begin{cases} 1, & \text{if } t = N, \\ (1 + (e^{\gamma\mu+\gamma^2\sigma^2/2}\alpha_{t+1}^{\gamma-1})^{\frac{1}{1-\gamma}})^{-1}, & \text{otherwise,} \end{cases} \tag{5.26}$$

and the value function is

$$V_t(X_t^\pi) = \frac{(X_t^\pi)^\gamma}{\gamma}(\alpha_t)^{\gamma-1}, \tag{5.27}$$

as derived in Appendix D. The problem does not have any control variables for allocation of wealth into the risky and riskless assets and is wealth independent, hence no heteroskedasticity will exist and Smearing Estimate (5.12) is accurate enough. Figure 5.1 shows the optimal consumption each time period for the different methods described in previous sections. A clear bias can be identified for the transformation without consideration to the re-transformation bias, indicating it will not give an accurate solution, while the two other methods based on Smearing Estimate (5.12) and Smearing Estimate with Controlled Heteroscedastisity (5.14) are very close to the true optimal value.

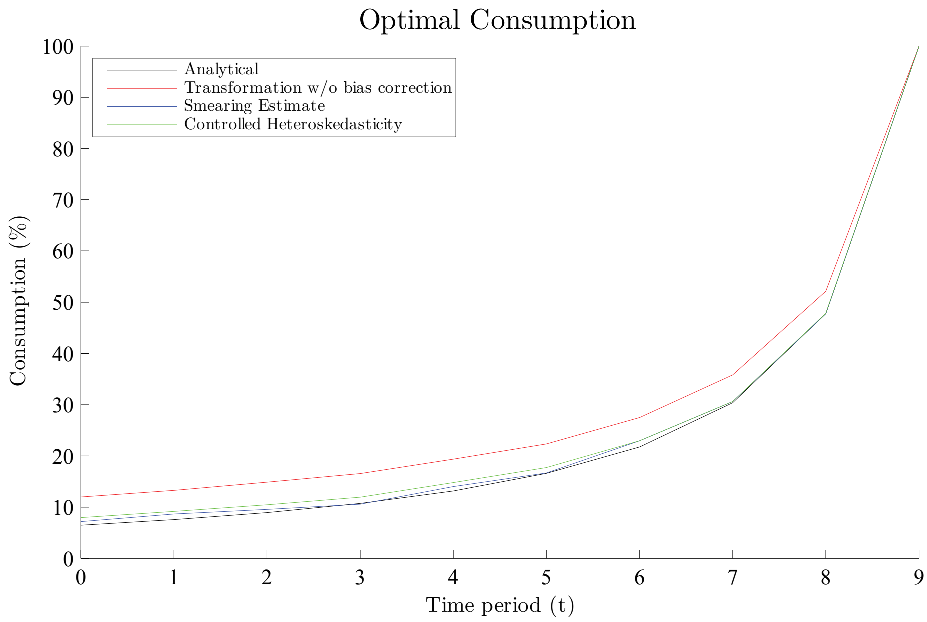

Figure 5.1: Optimal consumption α_t as a percentage proportion of wealth for four different solution methods.

5.5.2 Consumption and investment model

The second model we consider is based on the first, but extended with an additional control variable for risky asset allocation. The agent now consumes a proportion of wealth α_t and chooses to allocate a proportion $\delta_t \in [0, 1]$ of remaining wealth into a risky asset with stochastic return Z_t and the rest into a risk free asset with deterministic return $r = 0.03$. Hence, the decision variables are $\pi_t = (\alpha_t, \delta_t)$ and the transition function is

$$X_{t+1} = \mathcal{T}_t(X_t, \pi_t, Z_{t+1}) := X_t(1-\alpha_t)e^{\delta_t Z_{t+1} + (1-\delta_t)r}. \tag{5.28}$$

This transition function is approximately the same as the correct transition function for the specified allocation problem when returns are small. The closed-form solution (see Appendix D) gives optimal consumption

$$\alpha_t = \begin{cases} 1, & \text{if } t = N, \\ (1 + (e^{\gamma\delta_t\mu + \gamma^2\delta_t^2\sigma^2/2 + (1-\delta_t)\gamma r}\alpha_{t+1}^{\gamma-1})^{\frac{1}{1-\gamma}})^{-1}, & \text{otherwise,} \end{cases} \tag{5.29}$$

optimal risky allocation

$$\delta_t = \frac{r-\mu}{\gamma\sigma^2}, \tag{5.30}$$

and the value function is the same as in (5.27).

These changes introduce heteroskedasticity with respect to the control variable by allowing control of the disturbance term, and the Smearing Estimate no longer gives a correct solution. Figure 5.2 shows the optimal consumption for each time period and different methods, and Figure 5.3 shows the optimal risky allocation. Since the effect of the disturbance control is not transformed back, the risk is underestimated and therefore suggests full risky investment allocation as risk does not increase with higher allocation. The Smearing Estimate then includes a slight bias with regards to optimal consumption due to a constant risk being transformed, which is higher than the true optimal consumption but still underestimated in regards to δ_t. The transformation without bias correction, however, now seems to be accurate, but does in fact include two biases that happen to cancel each other out in this example. As the risk is ignored, the risk from the transformation without bias correction is underestimated and full risky investment is suggested. This in turn affects the expected capital growth, and together with the bias in the consumption decision the overall bias is almost cancelled out. Controlled Heteroskedasticity, on the other hand, considers the effect of disturbance terms once re-transformed and leads to an unbiased solution (but still subject to noise, regression and numerical errors).

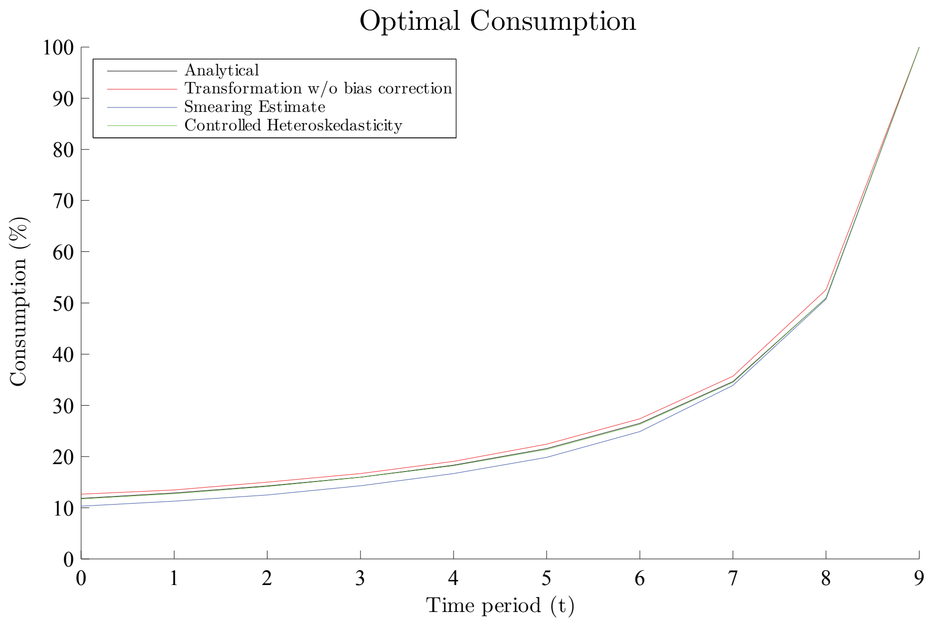

Figure 5.2: Optimal consumption α_t as a percentage proportion of wealth when the model allows risky investments, for four different solution methods.

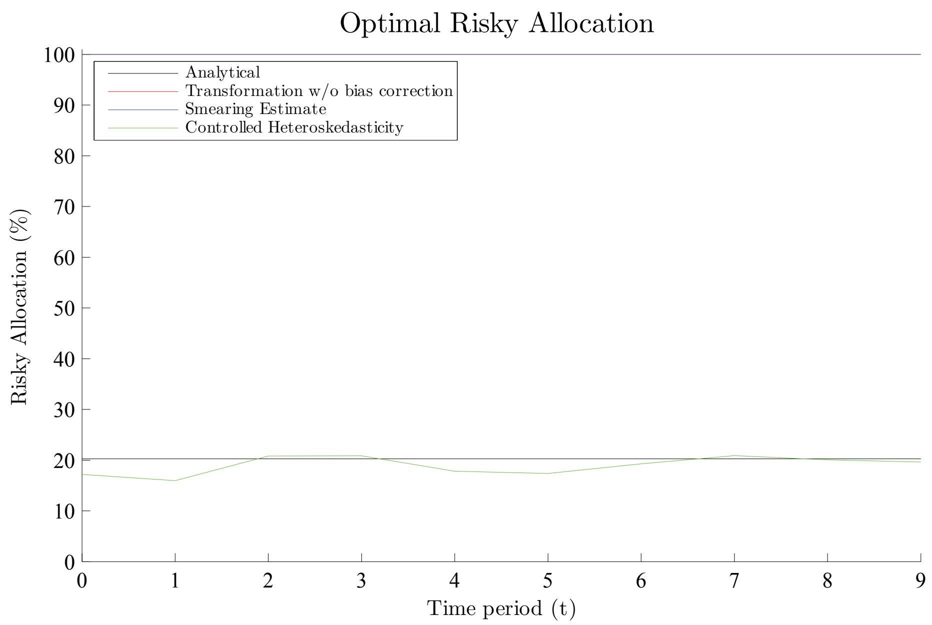

Figure 5.3: Optimal allocation of risky assets δ_t for four different solution methods.

5.5.3 Bounded solutions

Even though the accuracy of the estimated optimal control variables compared with the true optimal control is satisfactory, we performed a further analysis to see whether the calculated upper and lower bounds of the LSMC value function spans over the exact value function. The bounds are estimated as described in Section 5.4.3, in the case of Model 2 with Controlled Heteroskedasticity (5.14), considered in Section 5.2. The problem is solved with 20 independent iterations where each iteration involves $M = 10000$ independent sample paths. In order to estimate the solution and standard errors, the solutions of the 20 iterations are averaged and standard errors are calculated. Note that a utility function will always bias errors downwards, assuming it is concave, and the extreme curvature will quickly affect any bias. Therefore, the value function is compared on the transformed scale. Table 5.3 shows the lower bound Q_L (which is the value from using Algorithm 3 or 4), the upper bound Q_U from replacing the realised disturbance with the expectated value of the disturbance term. The absolute difference of the true optimal control parameters and the numerical solution for consumption, $|\Delta\alpha|$, and risky asset allocation, $|\Delta\delta|$, are shown as the average and the maximum difference. Regression is based on ordinary polynomials up to the n-th degree in both X_t and π_t, and results in the table are presented for $n = 2, 3, 4, 6, 8$.

Table 5.3: Bounded solutions and differences in control variables with different basis functions. The analytical solution of the problem is $H^{-1}(V_0) = 9119$.

nth degree	Q_{L}	Q_{U}	$\lvert\Delta\alpha\rvert_{\mathrm{avg}}$	$\lvert\Delta\alpha\rvert_{\mathrm{max}}$	$\lvert\Delta\delta\rvert_{\mathrm{avg}}$	$\lvert\Delta\delta\rvert_{\mathrm{max}}$
2	8735.9 (1.4)	8968.2 (2.0)	0.0011	0.0016	0.0285	0.0640
3	8742.9 (1.2)	9009.1 (6.4)	0.0011	0.0017	0.0225	0.0562
4	8745.7 (1.7)	9027.0 (4.6)	0.0010	0.0016	0.0215	0.0737
6	8742.5 (2.0)	9006.4 (7.4)	0.0012	0.0018	0.0255	0.0832
8	8739.3 (2.1)	9001.0 (7.2)	0.0014	0.0032	0.0235	0.0602

The difference between the bounds and the true value is due to the regression bias, and that an approximate model always will be suboptimal. As the upper bound given in Aïd et al. (2014) contains a significant amount of regression bias over time, which is further inflated by utility based objective functions, this value turns out to be *lower* than Q_{L} and has been omitted from the table. Using an upper bound based on the expected change of the stochastic variables will therefore reflect a more realistic upper bound, although it can still be less than the true solution due to regression and numerical errors or early stopping criteria in the optimisation of optimal control.

The difference between the approximation and the analytical solution equals to as little difference as 4% using the forth order polynomials. The optimal control tends to be close to the true optimal value. This can be seen when polynomials with a higher order than four are used, as the regression model now contains redundant predictors. The optimal control parameters are still within a valid range from their true value, even though the non-transformed value function starts to deviate quickly.

5.6 Conclusion

LSMC provides many advantages in dynamic programming. Firstly, it does not suffer from the 'curse of dimensionality' in the same way as other methods, and is therefore faster than numerical methods such as partial differential equation or quadrature based ones. Secondly, it does not impose restrictions on the stochastic variable dynamics, hence, even an empirical distribution is sufficient. Finally, it returns a parametric estimate in a feedback form of control which voids the need for a grid for control. However, there are many difficulties as well. It is an approximate method only and can be computationally intensive – especially in finding the optimal control variables for each sample. The basis functions are often difficult to find and highly problem-specific, and if they are not defined properly, substantial errors can pile up over multiple periods.

In this chapter the LSMC method was applied to stochastic control problems characterised by utility functions. We found that standard LSMC does not work well for these problems and suggest to perform regression on transformed value function and then accounting for the re-transformation bias. The bias correction function can be constructed in various ways depending on the type of problem. The Smearing Estimate can improve the accuracy of simpler problems without heteroskedasticity and control of disturbance, while more complex problems require Smearing Estimate with Controlled Heteroskedasticity if the heteroskedasticity depends on the state or control variables. The latter requires performing two regressions, but since the computational burden is on the optimisation and not the regression the additional computational cost is minimal. We further observed that the standard forward simulation stage of LSMC should be modified to achieve accurate results. In particular, we suggest to re-sample state variables independently at each time step to achieve better exploration of the state space. This occurs when the sample paths are simulated with control randomisation and the control has a significant influence on the transition of the state variable, thus all sample paths tend to end up in a small sub-domain of the state after simulation. By re-sampling the state variables each time step, we can ensure that the samples exist in the full state variable domain.

Chapter 6

Extension of retirement model with annuities and flexible housing decisions

6.1 Introduction

Australian pension policy is characterised by the means-test, which provides Age Pension payments for retirees with less wealth and/or income. The means-test raises a number of questions regarding optimal behaviour, such as optimal behaviour with respect to current or planned policies, but also regarding the validity of traditional knowledge in retirement modelling. One such insight is the 'fact' that a risk averse retiree tends to be better off by annuitising part of his/her wealth (Yaari, 1965; Davidoff et al., 2005; Milevsky and Young, 2007). A lifetime annuity is a financial product that pays a guaranteed income and insures against outliving one's assets (longevity risk). By purchasing an annuity the retiree gives up wealth that could potentially earn a higher return and which could be used as bequest. Even after the mortality credit[1], the payout rate is generally below the risk premium early in retirement, but insures the retirees from outliving their incomes. Risk averse agents[2], however, discount the risk premium more and value a protected income over potentially higher future consumption, thus annuitising more wealth (Iskhakov

[1]Mortality credit refers to the discounting from mortality risk of future income streams.

[2]This is true for rational investors only. Irrational investors, however, may value their current level of consumption too much and therefore defer annuitisation (Marín-Solano and Navas, 2010).

et al., 2015). There are alternative annuities that address the negative aspects of a lifetime annuity, such as Guaranteed Minimum Withdrawal Benefit and Guaranteed Minimum Death Benefit, which allow for equity growth and bequest motives respectively (Luo and Shevchenko, 2015). These products tend to be more expensive due to the additional benefits. The retiree therefore needs to find a balance between a guaranteed consumption and the possibility to leave bequest. Yaari (1965) showed that if no bequest motive is present, then full annuitisation is optimal. If such a bequest motive exists, however, annuitisation is still optimal but typically only partial (Davidoff et al., 2005; Friedman and Warshawsky, 1990), which is also the case when a certain consumption floor is present. Despite this, very few Australians annuitise any wealth (Iskhakov et al., 2015; Kingston and Thorp, 2005), which is consistent with retirees who receive other stable income streams (Inkmann et al., 2011; Dushi and Webb, 2004; Kingston and Thorp, 2005), with the exception of Switzerland where the majority of retirees do annuitise (Avanzi, 2010; Avanzi and Purcal, 2014). As the means-tested Age Pension provides an income stream that covers the calibrated consumption floor from Chapter 4, the Age Pension becomes a possible substitute for voluntary annuitisation.

Another important aspect of the means-test is the exclusion of any house asset, which opens up the possibilities for financial planning to maximise utility in retirement. Most Australian households do not convert housing assets into liquid assets in order to cover expenses in retirement, with the exception of certain events (such as the death of a spouse, divorce, or moving to an aged care facility) (Olsberg and Winters, 2005; Asher et al., 2017). However, by allocating more assets to the family home, the means-tested assets can be lowered which in turn results in more Age Pension received, and home equity can be retrieved later in retirement if needed. As with annuities, this raises the question whether retirees should access home equity, either by selling the home or through home equity products, or if the means-test crowds out such products as well. Sun et al. (2008) find that reverse mortgages are a very risky asset, owing to the uncertainty of the interest rates and housing markets. However, the decision to access home equity cannot be made purely for financial reasons and needs to be set in the context of typical Australian retirement behaviour. Due to both financial benefits and attachment to their home, and especially neighbourhood, retirees tend to stay homeowners late in life (Olsberg and Winters, 2005). The possibility to borrow money decreases with age, mainly due to having no labour income, and the retiree becomes increasingly locked into their home equity (Nakajima, 2017). An increasingly popular solution is therefore a reverse mortgage, which allows the retiree to borrow against home equity, up to a certain loan-to-value ratio (LVR) threshold. The LVR threshold tends to increase with age. The initial principal limit generally starts with 20-25% at age 65 (subject to expected interest rate and property value), which translates to either the lump sum or the present value of future payments, and increases 1% per year. The house equity is used as collateral and allows the retiree to access housing equity while maintaining residence in the house. The retiree can typically choose between six repayment options: lump sum, line of credit (allowing flexible amounts and payment times), tenure (equal monthly payments), term (tenure but with a fixed time horizon) and combinations of line of credit with either tenure or term (Chen et al., 2010). The loan is charged with either fixed or variable interest, but instead of requiring amortisation or interest payments they accumulate (although the retiree is free to make repayments at any time to reduce debt). The main benefits of such an arrangement are that it limits the risk as the loan repayments are capped at the house value, and allows the retiree to access more equity with age (contrary to traditional loans). However, interest rates are higher due to lending margin and insurance. Chiang and Tsai (2016) find that the desire for reverse mortgages is negatively correlated with the costs (application costs and insurance/spread added to the interest rate) as well as the income for a retiree, and according to Nakajima (2017) the loans are very expensive for retirees. In addition, if a lump sum is received and allocated to what is considered an asset in the means-test, such as a risky investment or simply a bank account, it will affect the Age Pension. On the other hand, if the funds are spent right away they will not have an impact on the Age Pension received.

Previous research found that the Age Pension crowds out decisions that otherwise are optimal (Iskhakov et al., 2015; Bütler et al., 2016). Here, it is evaluated whether such findings are consistent in a more realistic framework. Asher et al. (2017) finds evidence that few households use financial products to access home equity, such as reverse mortgages. For these reasons, we investigate whether the retiree is better off based on two additional control variables: borrowing against housing assets with a reverse mortgage, or simply up/downsizing housing in retirement. Since housing assets are exempt from the means-test, it might be optimal to over allocate in housing and then draw it down by reverse mortgage. Extending the model with more flexible decisions for homeowners is highly timely: in Australian Government (2017), the government announced that retirees will be able to deposit non-concessional contributions from the proceeds of selling their home into their super fund account (subject to additional conditions being met). The deposit is capped at $300,000 per retiree, hence couples can deposit twice the amount. The reason is to encourage downsizing in retirement, where the additional living space is no longer needed. As these rules will be in effect from the 1st of July 2018 they are not explicitly modelled in this chapter, but signifies the importance of understanding the effect that house equity related decisions has on the retiree.

The chapter is structured as follows. First, the 'benchmark' model is defined in Section 6.2 which is the foundation used in this chapter. In Section 6.3, additional optimal control with respect to annuitisation decisions and home equity access is modelled individually. The results of each extended model are evaluated in Section 6.4. Finally, the chapter is concluded in Section 6.5.

6.2 Benchmark model

The foundations for this chapter are based on the model presented in Chapter 3, with the same utility functions and parameters, but extend the model in several important aspects[3]. First, a stochastic risk-free rate is introduced. Second, a deposit account is now available in addition to the Allocated Pension account. Although the definition of a deposit account normally is that it only pays interest and has restrictions on withdrawals, we use this in a wider sense that allows financial investments, interest rate investments and yearly withdrawals and deposits with no restrictions on size.

[3]It should be noted that since the model was calibrated to data based on certain assumptions of deterministic variables, changing these to stochastic might have implications on the utility parameters. Using the same parameters does, however, function as a benchmark to evaluate the benefits of additional decisions and extensions to the model. We do not say in this case that the average Australian retiree is recommended to act based on the model solution, only that such a solution can show whether the retiree is better or worse off with regards to the decisions.

6.2.1 Additional dynamics and states

The stochastic risk-free rate is modelled as a Vasicek process

$$dr(t) = b(\bar{r} - r(t))dt + \sigma_R dB(t), \tag{6.1}$$

where $b \in (0,1]$ is the reversion to the mean, $\bar{r} \in \mathbb{R}^+$ is the mean the process reverts to, $\sigma_R \in \mathbb{R}^+$ is the volatility and $B(t)$ the standard Brownian motion. The corresponding process is discretised with yearly time steps, hence the process is

$$r_{t+1} = \bar{r} + e^{-b}(r_t - \bar{r}) + \sqrt{\frac{\sigma_R^2}{2b}(1 - e^{-2b})}\epsilon_{t+1}, \tag{6.2}$$

where $\epsilon \sim \mathcal{N}(0,1)$ is an i.i.d disturbance term. The Vasicek process is chosen as it allows for negative interest rates, which is suitable as the rate is defined in real terms. A negative interest rate would then indicate that inflation is higher than the nominal risk-free rate.

The deposit account $\widetilde{W}_t \in \mathbb{R}^+$ is an account which holds liquid wealth separate from the pension account, where the balance can be invested each period $[t, t+1)$ and is included in the means-test. It is assumed that the deposit account is invested in the same way as the Allocated Pension account, but the deposit account must pay taxes on any capital gains. The purpose of such an account is that the retiree will be able to save part of the Age Pension and/or drawdowns from the pension account when minimum withdrawals are larger than what is optimal to consume. Such an account is necessary later on in Section 6.3.1 and 6.3.2, as it is possible to receive lump sums but pension accounts do not allow funds to be added to them after retirement. The consumption each period consist of

$$C_t = \alpha_t(W_t + \widetilde{W}_t) + P_t, \tag{6.3}$$

where α_t determines the total drawdown from the deposit and pension account. It can be argued whether a second control variable for consumption from the deposit account is required, as the retiree now can choose what account to withdraw from (as long as the minimum withdrawal requirement in the pension account is satisfied). However, it is assumed the retiree first draws wealth from the pension account up to the minimum withdrawal rate ν_t each period, and in case optimal consumption exceeds this amount the difference is taken from the deposit account (as long as sufficient funds are available in the deposit account). Due to the deposit account attracting a tax on capital gains while the Allocated Pension account is tax-free, it will always be optimal to deplete the deposit account first, hence the assumption above. Transitions for the pension account and the deposit account depend on whether the deposit account can cover any desired drawdowns above the minimum withdrawal rates. Thus, if $\widetilde{W}_t(1-\alpha_t) > W_t(\alpha_t - \nu)$, the evolution for the pension account is

$$W_{t+1} = W_t(1-\nu_t) \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}). \tag{6.4}$$

For the deposit account, the evolution is

$$\begin{aligned}\widetilde{W}_{t+1} =& [\widetilde{W}_t(1-\alpha_t) + W_t(\nu - \alpha_t)] \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}) \\ &- \Theta(\widetilde{W}_t(1-\alpha_t) + W_t(\nu-\alpha_t), \delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}),\end{aligned} \tag{6.5}$$

where the function Θ calculates the tax on asset growth, and is defined as

$$\Theta(w,z) = 0.15w\max(z-1,0). \tag{6.6}$$

Note that only the deposit account attracts a tax on capital gains. For simplicity, it is assumed that any gains are realised each year, and that the tax rate is 15%[4]. If consumption is less than minimum withdrawals, the excess funds are stored in the deposit account. On the other hand, if $\widetilde{W}_t(1-\alpha_t) \leq W_t(\alpha_t-\nu)$, the deposit account is depleted ($\widetilde{W}_{t+1} = 0$) and the excess consumption comes from the pension account which evolves as

$$W_{t+1} = (W_t + \widetilde{W}_t)(1-\alpha_t) \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}). \tag{6.7}$$

[4]Due to the many tax offsets, rebates and investment options in retirement which can alter the effective tax rate, the tax rate has been set to a fixed 15% which equals the earnings tax on Self-Managed Super Funds.

6.2.2 Stochastic control problem definition

For the purpose of a complete definition of the benchmark model, it is defined in the stochastic control problem framework. For details regarding assumptions and further explanations please refer to Chapter 3.

- Denote a state vector as $X_t = (W_t, \widetilde{W}_t, G_t, H_t, r_t) \in \mathcal{W}\times\mathcal{W}\times\mathcal{G}\times\mathcal{H}\times\mathcal{R}$, where $W_t \in \mathcal{W} = \mathbb{R}^+$ and $\widetilde{W}_t \in \mathcal{W} = \mathbb{R}^+$ denotes the current level of liquid wealth in a pension account and a deposit account respectively. $G_t \in \mathcal{G} = \{\Delta, 0, 1, 2\}$ denotes whether the agent is dead, died in this period, is alive in a single household, or is alive in a couple household. The stages are sequential; hence, an agent that starts out as a couple becomes single when one spouse dies. $H_t \in \mathcal{H} = \mathbb{R}^+$ denotes the value of the home and $r_t \in \mathcal{R} = \mathbb{R}$ the stochastic real risk-free interest rate (thus can take on negative values).

- Denote an action space of $(\alpha_t, \delta_t, \varrho) \in \mathcal{A} = (-\infty, 1] \times [0, 1] \times \{0, [\frac{H_L}{\mathsf{W}}, 1]\}$ for $t = t_0$, and $(\alpha_t, \delta_t) \in \mathcal{A} = (-\infty, 1] \times [0, 1]$ for $t = t_0 + 1, ..., T - 1$. Here, $\varrho \in \{0, [\frac{H_L}{\mathsf{W}}, 1]\}$ is the proportion total wealth allocated to housing, $\alpha_t \in (-\infty, 1]$ denotes the proportion of wealth consumed and $\delta_t \in [0, 1]$ is the percentage of wealth allocated in the risky asset. The upper boundary of 1 indicates that the drawdown cannot be larger than the total wealth, nor invest more than 100% in risky assets; hence, borrowing is not allowed. However, negative values for drawdown are allowed as they represent savings from Age Pension payments into the deposit account. Housing requires a certain minimum down payment H_L, and cannot exceed total wealth at retirement.

- Denote an admissible space of state-action combination as $D_t(x_t) = \{\pi_t(x_t) \in \mathcal{A} \mid \alpha_t \geq \frac{\bar{c}_d - P_t}{W_t + \widetilde{W}_t}\}$, which includes the possible actions for the current state and indicates that withdrawals must be sufficiently large to cover the necessary consumption floor.

- There exist transition functions for the state variables W_t, $\widetilde{W}_t$ and r_t. As housing is constant in retirement, $H_{t+1} = H_t$. Define the total transition function

$$T_t(W_t, \widetilde{W}_t, r_t, \alpha_t, \delta_t, z_{t+1}) = \begin{bmatrix} T_t^W(W_t, \alpha_t, \delta_t, z_{t+1}, r_{t+1}) \\ T_t^{\widetilde{W}}(\widetilde{W}_t, \alpha_t, \delta_t, z_{t+1}, r_{t+1}) \\ T_t^r(r_t) \end{bmatrix}. \tag{6.8}$$

Here, $T_t^W(\cdot)$ is the transition function for the pension account

$$T_t^W(\cdot) := W_{t+1} = \begin{cases} W_t(1-\alpha_t) \\ \quad \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \text{if } \widetilde{W}_t(1-\alpha_t) > W_t(\alpha_t - \nu), \\ (W_t(\nu-\alpha_t) + \widetilde{W}_t(1-\alpha_t)) \\ \quad \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \text{otherwise,} \end{cases} \tag{6.9}$$

where z_{t+1} and r_{t+1} is the realisation of the return on the stochastic investment portfolio and risk-free rate respectively, over $(t, t+1]$. We assume that the agent is small and cannot influence the asset price. $T_t^{\widetilde{W}}(\cdot)$ is the transition function for the deposit account

$$T_t^{\widetilde{W}}(\cdot) := \widetilde{W}_{t+1} = \begin{cases} (\widetilde{W}_t(1-\alpha_t) + W_t(\nu-\alpha_t)) \\ \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}) \\ - \Theta(\widetilde{W}_t(1-\alpha_t) + W_t(\nu-\alpha_t), & \text{if } \widetilde{W}_t(1-\alpha_t) > W_t(\alpha_t-\nu), \\ \delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), \\ 0, & \text{otherwise.} \end{cases} \tag{6.10}$$

Finally, $T_t^r(\cdot)$ is the transition function for the stochastic risk-free rate, which is based on equation (6.2), hence

$$T_t^r(r_t) := r_{t+1} = \bar{r} + e^{-b}(r_t - \bar{r}) + \sqrt{\frac{\sigma_R^2}{2b}(1-e^{-2b})}\epsilon_{t+1}. \tag{6.11}$$

- Denote the stochastic transitional kernel as $Q_t(dx'|x, \pi_t(x))$, which represents the probability of reaching a state in $dx' = (dw_{t+1}, d\widetilde{w}_{t+1}, g_{t+1}, dr_{t+1})$ at time

$t+1$ if action $\pi_t(x)$ is applied in state x at time t. Since the transition function is based on the stochastic risky return Z_{t+1} and the stochastic risk-free rate r_{t+1}, which are Markovian, the transition probability for W_{t+1}, $\widetilde{W}_{t+1}$ and R_{t+1} are determined by the distributions, where $Z_{t+1} \overset{i.i.d}{\sim} \mathcal{N}(\mu - \widetilde{r}, \sigma_Z^2)$ and $r_{t+1} \overset{i.i.d}{\sim} \mathcal{N}(\bar{r} + e^{-b}(r_t - \bar{r}), \frac{\sigma_R^2}{2b}(1 - e^{-2b}))$. As the problem is solved with a simulation based method, the stochastic kernel with respect to the financial stochastic variables does not have to be explicitly defined. The survival probabilities will, however, be implemented directly in the calculations. Let $q(g_{t+1}, g_t)$ denote $\Pr[G_{t+1} = g_{t+1} \mid G_t = g_t]$. The stochastic kernel is then given by

$$\begin{aligned}
&Q_t(dx'|x, \pi_t(x)) \\
&\quad = \Pr[W_{t+1} \in dw_{t+1}, \widetilde{W}_{t+1} \in d\widetilde{w}_{t+1}, G_{t+1} = g_{t+1}, r_{t+1} \in dr_{t+1} \mid X_t = x_t] \\
&\quad = \Pr[W_{t+1} \in dw_{t+1}, \widetilde{W}_{t+1} \in d\widetilde{w}_{t+1}, r_{t+1} \in dr_{t+1} \mid W_t = w_t, \widetilde{W}_t = \widetilde{w}_t, r_t] \\
&\qquad \times q(g_{t+1}, g_t).
\end{aligned} \tag{6.12}$$

The probabilities for family status are defined as

$$\begin{aligned}
&q(2,2) = p_t^{\mathrm{C}}, \quad q(1,2) = 1 - p_t^{\mathrm{C}}, \\
&q(1,1) = p_t^{\mathrm{S}}, \quad q(0,1) = 1 - p_t^{\mathrm{S}}, \\
&q(\Delta, 0) = q(\Delta, \Delta) = 1,
\end{aligned} \tag{6.13}$$

where p_t^{C} is the probability of surviving for one more year as a couple or p_t^{S} as a single. All other transition probabilities for family status are 0.

- The reward function depends on the G_t state as defined in equation (6.14). If the agent is alive, he/she receives a reward based on consumption. Thus, if the agent is alive, he/she receives a reward based on consumption U_C and housing U_H. If he/she died during the year, the reward comes from the bequest U_B, and if he/she is already dead, there is no reward. Note that the reward received when the agent is alive depends on whether the state is a couple or single household owing to differing utility parameters and Age Pension thresholds.

$$R_t(W_t, \widetilde{W}_t, G_t, \alpha_t, H_t) = \begin{cases} U_C(C_t, G_t, t) + U_H(H_t, G_t), & \text{if } G_t = 1, 2, \\ U_B(W_t + \widetilde{W}_t + H_t), & \text{if } G_t = 0, \\ 0, & \text{if } G_t = \Delta. \end{cases} \tag{6.14}$$

The utility function for consumption is defined as

$$U_C(C_t, G_t, t) = \frac{1}{\psi^{t-t_0}\gamma_d}\left(\frac{C_t - \overline{c}_d}{\zeta_d}\right)^{\gamma_d}, \quad d = \begin{cases} \text{C}, & \text{if } G_t = 2 \quad \text{(couple)}, \\ \text{S}, & \text{if } G_t = 1 \quad \text{(single)}, \end{cases} \tag{6.15}$$

where $\gamma_d \in (-\infty, 0)$ denotes the risk aversion, $\overline{c}_d$ is the consumption floor, ζ_d is a scaling factor that normalises utility between couple and single household. These parameters are subject to family state G_t. Finally, $\psi \in [1, \infty)$ is a "health" proxy to control for decreasing consumption with age.

The bequest function is defined as

$$U_B(W_t + \widetilde{W}_t + H_t) = \left(\frac{\theta}{1-\theta}\right)^{1-\gamma_\text{S}} \frac{\left(\frac{\theta}{1-\theta}a + W_t + \widetilde{W}_t + H_t\right)^{\gamma_\text{S}}}{\gamma_\text{S}}, \tag{6.16}$$

where W_t denotes the liquid assets available for bequest, γ_S denotes the risk aversion parameters of a singles household, $\theta \in [0, 1)$ the utility parameter for bequest preferences over consumption, and $a \in \mathbb{R}^+$ the threshold for luxury bequest.

Finally, the utility function for housing is defined as

$$U_H(H_t, G_t) = \begin{cases} \frac{1}{\gamma_\text{H}}\left(\frac{\lambda_d H_t}{\zeta_d}\right)^{\gamma_\text{H}}, & \text{if } H_t > 0, \\ 0, & \text{if } H_t = 0, \end{cases} \tag{6.17}$$

where γ_H is the risk aversion parameter for housing (different from risk aversion for consumption and bequest), H_t is the value of the family home and $\lambda_d \in (0, 1]$ is the housing preference defined as a proportion of the market value.

- The terminal reward at $t = T$ is given by

$$\widetilde{R}(W_T, \widetilde{W}_T, G_T, H_T) = \begin{cases} U_B(W_T + \widetilde{W}_T, H_T), & \text{if } G_T \geq 0, \\ 0, & \text{if } G_T = \Delta. \end{cases} \tag{6.18}$$

- The discount factor β, which was previously set in relation to the deterministic interest rate, does not represent financial discounting but rather personal impatience hence remains fixed. The discount factor remains the same, with $\beta_{t,t+1} \in (0, 1]$.

In order to solve the model, the Least-Squares Monte Carlo approach in Chapter 5 will be used.

6.3 Extensions

The model is now extended to areas of interest: annuitisation (extension 1) and scaling housing/reverse mortgaging (extension 2). Note that extension 1 does not apply in extension 2 and vice versa - they are separate and independent extensions which isolate the impact each extension has on optimal control. Any parameterisation not explicitly stated remains the same as in Chapter 4, where the most recent policy and parameters are used for the Age Pension in Table 4.4.

6.3.1 Extension 1 - Annuitisation

The argument why the Australian market has shown such a lack of interest in annuities comes down to that the Age Pension is indirectly an indexed life annuity which pays a known and increasing amount as wealth and income decreases, hence crowding out annuity purchased (Iskhakov et al., 2015; Bütler et al., 2016). This provides an implicit insurance against both longevity and financial risk, which otherwise is the main argument to annuitise. If annuities were exempt of the means-test, then it would be reasonable to expect an increased interest in annuities. However, the annuity value as well as annuity payment is included in the asset-test. Any annuitisation would therefore give up 'free' money if the means-test is binding, as well as give up potential equity growth, unless the annuity is of any equity-linked type.

Annuity pricing

The focus is on a single household, since the joint mortality risk of a couple household increases the price of the annuity, thus making it less attractive to annuitise. The retiree can each year decide if he/she wants to annuitise any wealth, hence making the annuity indirectly a deferred one by saving wealth in retirement in order to annuitise later (similar to Milevsky and Young (2007)). This introduces the possibility for the retiree to receive additional equity growth on the wealth yet to be annuitised, although with the risk associated, but without requiring more complex annuity products.

Assume an immediate lifetime annuity that is fairly priced (hence there are no commercial markups or fees) with constant real payments, where the actuarial present value can be written as

$$a_t(y) := \sum_{i=t+1}^{T} {}_ip_t J(t,i,y), \tag{6.19}$$

where $J(t,i,y)$ represents the price of a zero coupon bond at time t with maturity i and face value y (the constant real annuity payment, hence adjusted for inflation), ${}_ip_t$ is the probability of surviving from year t to i. The price of this kind of annuity equals a portfolio of mortality risk weighted bonds with maturity times from $t+1$ up to T. The bond is priced with yearly time periods, hence the bond with maturity time t_M is

$$J(t,t_M,y) = y\mathbb{E}^{\widetilde{Q}}[e^{-\int_t^{t_M} r(\tau)d\tau}] := ye^{-r(t,t_M)(t_M-t)}, \tag{6.20}$$

where $\widetilde{Q}$ is the risk-neutral measure and $r(t,t_M)$ is the zero rate. The corresponding Vasicek risk-neutral process is

$$dr(t) = [b(\bar{r} - r(t)) - \lambda\sigma_R]dt + \sigma_R d\tilde{B}(t). \tag{6.21}$$

The formulas for the bond price and corresponding zero rate can easily be calculated (see, e.g., Hull (2012))

$$r(t,t_M) = \frac{-\ln A(t,t_M) + B(t,t_M)r_t}{t_M - t}, \tag{6.22}$$

where

$$A(t,t_M)=\exp\left[\frac{(B(t,t_M)-t_M+t)(b^2(\bar{r}-\lambda\sigma_R/b)-\sigma_R^2/2)}{b^2}-\frac{\sigma_R^2 B(t,t_M)^2}{4b}\right], \tag{6.23}$$

$$B(t,t_M)=\frac{1-e^{-b(t_M-t)}}{b}, \tag{6.24}$$

and λ is the market price of risk. Equation (6.22) gives the zero rate in a risk-neutral world, hence the full term structure can be determined. In order to fit the real risk-free rate parameters, which are needed to find the correct discounting of the annuity payment, the process outlined in Hull (2012) is used. First the risk-free rate r_t process needs to be parameterised from real data. The yearly Australian real deposit rate is chosen to represent a real risk-free rate which the retiree has access to, where the dataset contains real deposit rates from 1989–2017. Then parameters of the Vasicek model are estimated using Maximum Likelihood method applied to the discretized version of the Vasicek model (equation 6.2) as follows.

$$\max_{b,\bar{r},\sigma}\sum_{i=1}^{n}\left(-\frac{1}{2}\ln\left(\frac{\pi\sigma_R^2}{b}(1-e^{-2b})\right)-\frac{(r_i-\bar{r}-e^{-b}(r_{i-1}-\bar{r}))^2}{\frac{\sigma_R^2}{2b}(1-e^{-2b})}\right), \tag{6.25}$$

where r_i is the observed real deposit rate at time t_i. The calibrated parameters are $\hat{b}=0.55$, $\hat{\bar{r}}=0.024$ and $\hat{\sigma}_R=0.029$, which are estimates of the parameters in the real world. The present real risk-free[5] rate is set to $r_0=0.0045$. To convert them into risk-neutral parameters (in order to derive the term structure of interest rates), it is necessary to estimate λ by minimising the squared difference between the term structure and the zero coupon market rates[6]. The term structure is generated from the zero rate function in equation (6.22) using the process in equation (6.21). The market price of risk is estimated by minimizing the mean squared errors for each trading date $t_i, i=1,...,n$ with maturity $T_j, j=1,..,10$

$$\min_{\lambda}\sum_i\sum_j\left(r(t_i,t_i+T_j)-r_{i,j}^{obs}\right)^2, \tag{6.26}$$

where $r_{i,j}^{obs}$ represents the observed yield at date t_i with maturity T_j. The estimate

[5] 1 year deposit rate from the Commonwealth Bank in April 2017.

[6] Taken from `https://www.quandl.com/data/RBA/F17_0-Zero-Coupon-Interest-Rates-Analytical-Series-Yields`

comes out as $\lambda = -0.14$, hence the risk-neutral parameter for the mean rate is $\bar{r} = \hat{\bar{r}} - \lambda\hat{\sigma}_R/\hat{b} = 0.031$, and the other equals the estimates where $b = \hat{b} = 0.55$ and $\sigma_R = \hat{\sigma}_R = 0.029$. The present value of the annuity can now be calculated. Note that we only attempt to define a model and solution that is useful, and are not trying to predict the market. The parameterisation and pricing of annuities are for illustration of the solution only.

Problem definition

In the context of the life cycle model, the retiree can at any time $t_0, ..., T-1$ make a (non-reversible) decision $\iota_t \in [0, 1-\alpha_t]$ to annuitise a proportion of liquid wealth $(W_t + \widetilde{W}_t)$. As the annuity is of the type annuity-immediate, the annuity payment is received in the end of the period (hence equals the time at $t+1$), but the cost is discounted to t and affects the wealth immediately in order to protect future consumption. Any decision to annuitise adds to the new state variable $Y_t \in \mathcal{Y} = \mathbb{R}^+$, which holds the information of the size of annuity payments each period. The transition function for the state variable is

$$T_t^Y(Y_t, y) := Y_{t+1} = Y_t + y, \tag{6.27}$$

where y is found from equation (6.19) by setting the decision to annuitise equal to the actuarial present value, $a_t(y) = \iota_t(W_t + \widetilde{W}_t)$, and solving for y. Note that y will always be non-zero due to the non-reversibility of the annuitisation decision. The transition functions for W_t and $\widetilde{W}_t$ is updated to

$$T_t^W(\cdot) := \begin{cases} W_t(1-\alpha_t-\iota_t) & \\ \quad \times\,(\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \text{if } \widetilde{W}_t(1-\alpha_t-\iota_t) > W_t(\alpha_t+\iota_t-\nu), \\ (W_t+\widetilde{W}_t)(1-\alpha_t-\iota_t) & \\ \quad \times\,(\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \text{otherwise.} \end{cases} \tag{6.28}$$

$$T_t^{\widetilde{W}}(\cdot) := \begin{cases} (\widetilde{W}_t(1-\alpha_t-\iota_t) + W_t(\nu-\alpha_t-\iota_t)) & \\ \times (\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}) & \text{if } \widetilde{W}_t(1-\alpha_t-\iota_t) > \\ -\Theta(\widetilde{W}_t(1-\alpha_t-\iota_t) + W_t(\nu-\alpha_t-\iota_t), & W_t(\alpha_t+\iota_t-\nu), \\ \delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \\ 0, & \text{otherwise.} \end{cases} \tag{6.29}$$

Any annuitisation is reflected by increased consumption in equation (6.15), hence the input for the utility from consumption becomes $U_C(\alpha_t(W_t + \widetilde{W}_t) + P_t + Y_t, G_t, t)$ as consumption is based on not only wealth drawdown and Age Pension, but now also annuity payments. Annuities need to be handled differently in the means-test. Any annuities which are not account-based (which is the type modelled here) are not included as deemed assets in the income-test, and are assessed based on the income they provide with a deduction for part of the annuity value (Department of Social Services, 2016). The definition of annuity income for the income-test is

$$y - \frac{a_{t_x}(y) - a_T(y)}{T - t_x}, \tag{6.30}$$

where t_x is the annuity purchasing time and $a_T(y)$ represents any residual capital value (which always will be zero as per the annuity definition). Similarly, the assessment of the annuity in the asset-test differs from other financial assets. The purchase price of the annuity is used, and a deduction to represent the decrease in value for each year is applied

$$a_{t_x}(y) - \frac{a_{t_x}(y) - a_T(y)}{T - t_x}(t - t_x). \tag{6.31}$$

These rules cause some implications to the model, as it will require additional state variables in terms of annuity purchase price and annuity purchasing time (which complicates the problem definition further as it is allowed to add on to annuities later in retirement). To avoid this, the calculations in equations (6.30) and (6.31) are approximated by using equation (6.19) to re-value the annuity at the current time given the known annuity payment. The assessment part for the income-test becomes approximately $y - \frac{a_t(y)}{T-t}$ and the asset-test assessment approximately equal to $a_t(y)$. In Figure 6.1 it can be seen that these methods are approximately equal, which voids the two additional dimensions in the model. Even if a solution using LSMC technically could handle the additional states, it is preferred to avoid this as the additional state variables will have a very minor impact on the value function but are prone to unnecessary regression errors. The means-test pension functions need to be updated though. The function for the income-test becomes

$$P_{\mathrm{I}} := P^d_{\max} - \left(P_{\mathrm{D}}(W_t) + y - \frac{a_t(y)}{T-t} - L^d_{\mathrm{I}} \right) \varpi^d_{\mathrm{I}}, \tag{6.32}$$

and the function for the asset-test

$$P_{\mathrm{A}} := P^d_{\max} - \left(W_t + a_t(y) - L^{d,h}_{\mathrm{A}} \right) \varpi^d_{\mathrm{A}}. \tag{6.33}$$

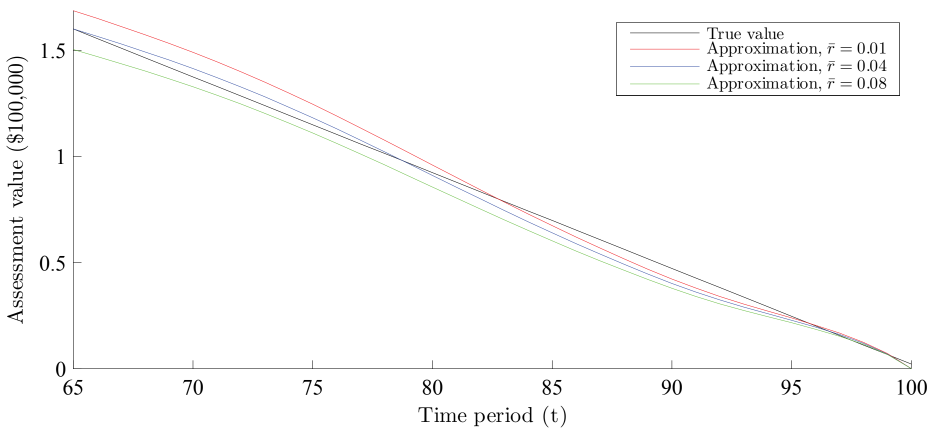

Figure 6.1: Comparison between the true value of the annuity assessment for the asset-test, compared with the approximation under three different interest rate scenarios ($\bar{r}_0 = 0.01\%$, $\bar{r}_0 = 0.04\%$ and $\bar{r}_0 = 0.08\%$). The annuity is priced based on a $10,000 annuity payment. The middle scenario reflects the calibrated risk-free rate.

Finally, the extended model requires some additional constraints which are reflected in the admissible action space. The lower bound of consumption drawdown now contains any annuity payments, which the retiree can choose to save in the deposit account instead of consuming them. The total drawdown (drawdown for consumption and for allocation to annuity purchase) cannot exceed total wealth, although allocation to annuity purchases can exceed total wealth if the retiree decides to save part of the Age Pension and annuity payment. The admissible action space is therefore updated to

$$D_t(x_t) = \left\{ \pi_t(x_t) \in \mathcal{A} \;\middle|\; \alpha_t \geq \frac{\bar{c}_d - P_t - Y_t}{W_t + \widetilde{W}_t},\ \alpha_t + \iota_t \leq 1 \right\}. \tag{6.34}$$

6.3.2 Extension 2 - Scaling housing and reverse mortgages

The second main extension to the model allows the retiree to either scale the housing by selling the current home and acquiring a new one of a different size or standard. Although downsizing is more common in retirement, especially in the case of a spouse passing away (Olsberg and Winters, 2005; Asher et al., 2017), the retiree is allowed to both up- and downscale at any point in time by making a decision $\tau_t \in [-1, \infty)$. A positive value represents the proportion of the current house value added to housing (upsizing from the current house), where the transition function for housing becomes

$$T_t^H := H_{t+1} = H_t(1 + \tau_t). \tag{6.35}$$

The decision variable is therefore bounded below by the current house value, and the upper bound depends on wealth. Decision is made in the start of each period and any house scaling is assumed to be instantaneous (no delay between the decision, the sale of the house and buying a new on). To capture the illiquid nature of housing assets, a proportional transaction cost ø applies. This will reflect actual costs associated with a sale of the house, as well as avoiding the risk of the optimal decision being a gradual yearly change in the housing asset. The transaction cost only affects the sale of the house, as any transaction cost for a new purchase is assumed to be absorbed by the other party. However, taxes on asset appreciation of the home are not modelled. Although it is likely to have an impact on the decision in reality, most retirees who are not homeowners at the time of retirement do not purchase a home (Nakajima, 2017), hence, the retirees who change houses will already be homeowners. Due to the housing boom that has been present in Australia since the 1980s, many homes will have a substantial proportion of unrealised asset gains if the home has been kept for a long time, but a much smaller portion if the home was purchased recently. Due to this variation, taxes on house asset gains have been omitted with the expectation that optimal decisions will be an upper bound for how often scaling of homes occurs, or for the size of the scaling.

The retiree can also choose to take out a reverse mortgage on the house. The assumptions of the loan structure is based on Shao et al. (2015), although not limited to a single payment at issuance. Define $L_t \in \mathcal{L} = \mathbb{R}^+$ to be the loan value at time t. The retiree can at any time make the decision to loan a certain proportion $l_t \in [0, \widetilde{L}_t]$ of the house value up to the threshold $\widetilde{L}_t$. The loan is based on a variable interest rate, where the outstanding loan amount accumulates over time. It is possible to increase an existing loan at any time up to $\widetilde{L}_t$, which is given by

$$\widetilde{L}_t = H_t I(t) \tag{6.36}$$

where $I(t)$ is a function for the principal limit (maximum LVR ratio) which changes with age, and is defined as

$$I(t) = 0.2 + 0.01(\min(85, t) - 65). \tag{6.37}$$

The maximum LVR therefore starts at 20% for age 65, which increases with 1% per year to a maximum of 40% at age 85[7]. The retiree is not liable to repay part of the loan in the case where the loan value exceeds the LVR or the house value due to accumulated interest (cross-over risk). If the retiree dies, or decides to sell the house, any remaining house value after loan repayments goes to wealth (and can be bequeathed). As Australian reverse mortgages include a ‘no negative equity guarantee’[8], the retiree (or the beneficiaries) are not required to cover any remaining negative house asset if $L_t > H_t$ at time of death or if the house is sold[9]. From the lender’s point of view, this results in two main risks: house price risk and longevity risk. If the house price decreases, or the retiree lives too long so that the loan value accumulates over the house value, the lender is liable for any losses unless these are forwarded to a third party via insurance. Increased interest rates can also speed up compounding of the loan, which increases crossover risk. These risks are in practice

[7] The parameterisation follows ‘Equity Unlock Loan for Seniors’ offered by the Commonwealth Bank, but does not impose a minimum or maximum dollar value for the loans.

[8] The guarantee is still subject to default clauses which can negate the guarantee, such as not maintaining the property, malicious damage to the property by the owner, failure to pay council rates and failure to inform the provider that another person is living in the house.

[9] Even if the possibility exists, it will not be optimal to sell the house if the net house asset is negative as the retiree will give up ‘free’ housing utility and receive no extra wealth. The exception is a significant upsizing at old age, which is not very likely.

covered with a mortgage insurance premium rate added to the loan, in addition to any lending margin required by the lender. The loan-value state therefore requires a transition function, and evolves as

$$T_t^L := L_{t+1} = (L_t \mathbb{I}_{\tau_t=0} + l_t H_t(1+\tau_t))e^{r_{t+1}+\varphi} \tag{6.38}$$

where $\mathbb{I}_{\tau_t=0}$ is the indicator symbol if no changes to house assets are made, and φ represents the lending margin and mortgage insurance premium combined. In the case $\tau_t \neq 0$, any outstanding loan value must be repaid, hence the loan is reset and a new loan can be taken out subject to the new house value. The costs of any decision (house transaction cost (ø), the difference in house assets in case of scaling and repayment of loan) is reflected in the wealth process. Let

$$\Delta_t^H = \frac{l_t H_t(1+\tau_t) - \mathbb{I}_{\tau_t \neq 0}\left(H_t(\tau_t + ø) + L_t\right)}{W_t + \widetilde{W}_t} \tag{6.39}$$

represent all changes to wealth from house scaling and reverse mortgage decisions as a proportion of current wealth, where $\mathbb{I}_{\tau_t \neq 0}$ is the indicator symbol if any scaling of housing occurs. Then the transition functions for the wealth states can be defined as

$$T_t^W(\cdot) := \begin{cases} W_t(1-\alpha_t-\Delta_t^H) & \text{if } \widetilde{W}_t(1-\alpha_t-\Delta_t^H) > W_t(\alpha_t+\Delta_t^H-\nu), \\ \quad \times\,(\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \\ (W_t+\widetilde{W}_t)(1-\alpha_t-\Delta_t^H) & \text{otherwise}, \\ \quad \times\,(\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \end{cases} \tag{6.40}$$

$$T_t^{\widetilde{W}}(\cdot) := \begin{cases} (\widetilde{W}_t(1-\alpha_t-\Delta_t^H) + W_t(\nu-\alpha_t-\Delta_t^H)) & \\ \quad \times\,(\delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}) & \text{if } \widetilde{W}_t(1-\alpha_t-\Delta_t^H) > \\ \quad -\,\Theta(\widetilde{W}_t(1-\alpha_t-\Delta_t^H) + W_t(\nu-\alpha_t-\Delta_t^H), & W_t(\alpha_t+\Delta_t^H-\nu), \\ \delta_t e^{z_{t+1}} + (1-\delta_t)e^{r_{t+1}}), & \\ 0, & \text{otherwise.} \end{cases} \tag{6.41}$$

In addition to new transition functions, the bequest function needs to include the house asset after any reverse mortgage has been repaid, and becomes $U_B(W_t +$

$\widetilde{W}_t, \max(H_t - L_t, 0))$.

New constraints need to be imposed on the control variables. The option to take out (or add to) a reverse mortgage is bounded from above by the difference of any outstanding mortgage and the loan-to-value ratio $\widetilde{L}_t$, hence

$$l_t \leq \max\left(0, \widetilde{L}_t - \frac{L_t \mathbb{I}_{\tau=0}}{H_t(1+\tau_t)}\right). \tag{6.42}$$

Note that if the control variable for scaling housing is not 0, any outstanding reverse mortgage must be paid back in full and a new reverse mortgage is available against the new house value. The max-condition in the formula is to ensure that the upper bound does not fall below the lower bound to ensure a feasible solution. For the scaling of housing, an upper bound for how much the house asset can be increased is determined by available wealth after costs associated with selling the current house (and repaying any outstanding reverse mortgage) and allocating additional wealth to the new house

$$\tau_t \leq \frac{W_t + \widetilde{W}_t - \mathbb{I}_{\tau\neq 0}\left(\o H_t + L_t\right)}{H_t}. \tag{6.43}$$

The lower bound is simply -1, since the retiree cannot downscale further than selling the house and not buying a new one, and the cost associated with the sale is reflected in the transition functions for the state variables. Finally, drawdown still needs to cover any consumption that exceeds the Age Pension received, but no longer has an upper bound of 1 as the maximum amount possible to draw down depends on how housing decisions and new mortgages affect current wealth. This is fully covered in the budget constraint

$$\alpha_t(W_t + \widetilde{W}_t) + \mathbb{I}_{\tau\neq 0}\left(\o H_t + L_t\right) - l_t H_t(1+\tau_t) - (W_t + \widetilde{W}_t) \leq 0. \tag{6.44}$$

The constraint specifies the total effect control variables have on wealth, where it ensures that the wealth is enough to cover consumption and housing costs in the case of scaling housing (including repaying any outstanding reverse mortgage) and grows if an additional reverse mortgage is taken out.

As for parameterisation, the transaction cost of selling is set to $\o = 6\%$ as in Nakajima (2017) and Shao et al. (2015). The markup to the interest rate is set according to Chen et al. (2010), with $\varphi = 0.0242$, but does not require a starting cost to access the loan. In addition, it is assumed there is no current debt on the house (or that it is used as security for other liabilities), and that there are no monthly fees in addition to φ.

6.3.3 Numerical solution

Both models are solved with Algorithm 4 as outlined in Chapter 5. The approximation of the conditional value function is made with ordinary least-squares regression, where the basis function consists of fourth order ordinary polynomials of the state and control variables. The exception is for extension 2, where the state variable for the outstanding loan value L_t is replaced with the covariate $\max(0, H_t - L_t)$ as this is how it appears in the bequest function. To avoid the transformation bias in the algorithm, the combination of the Smearing Estimate and Controlled Heteroskedasticity defined in equation (5.14) is used. The solution is run with 10,000 sample paths, and the optimisation of the variables is performed with a grid search algorithm.

6.4 Results

6.4.1 Extension 1: Annuitisation

The optimal annuitisation is expected to differ from previous research due to a number of reasons. In both Iskhakov et al. (2015) and Bütler et al. (2016), the retirement is modelled with a starting wealth that is assumed to be fully consumed and cannot be bequeathed. This means that the level of annuitisation identified given a certain wealth, age and parameters is optimal on a relative basis compared to alternative investment options in order to maximise consumption each time period. Since the model in this thesis is calibrated to the behaviour of Australian retirees, where wealth appears in the bequest function, the annuitisation rate is expected to be lower. Similarly, as consumption declines with age, any desired consumption above the consumption floor which can be covered with annuitising early in retirement is not as desirable at older ages. In many cases this excess consumption is fully covered by the Age Pension payments. In addition to this, Iskhakov et al. (2015) do not allow for a risk-free rate, hence the annuity is the only (non-reversible) option to access risk-free investments. As the extension model allows the retiree to choose a risk-free allocation, this option can decrease the annuitisation further.

Figure 6.2 compares the results from the definition in Section 6.3.1 with the scenarios where no risk-free investment is available ($\delta_t = 1 \, \forall t$), and the scenario where the annuity value for each period is included in the bequest function. Each scenario assumes that no prior annuitisation has been done. The case where no risk-free asset is available is almost identical to the default case, indicating that retirees prefer annuities over risk-free investments due to the mortality credit, but the case with bequeathable annuities are significantly different. If the interest rates happen to be higher than normal, then allocation to annuities is slightly higher as well, even if the interest rate is expected to revert back to normal levels. The results show that it is still optimal with part annuitisation in a more realistic simulation environment, although at a low level. This level decreases with age, and becomes close to constant for higher levels of wealth. For low levels, where full Age Pension is received, the optimal allocation quickly goes towards zero. A retiree with \$500,000 in liquid wealth at retirement optimally allocates 22% to annuities, which results in approximately \$6,500 in annual annuity payments. If the decision is deferred to age 85, the optimal annuitisation is now lower at 12% for the same wealth, but the resulting annual payments are higher at \$16,500. Although an Australian retiree has a lower desire for consumption at an older age, the mortality credit at this age is significant and the retiree can access a large boost in yearly consumption for a relatively small wealth sacrifice, resulting in higher overall utility.

To set this in relation to previous research, the results from Iskhakov et al. (2015) suggest on average a higher level of annuitisation, where the range of the authors' different risk preference and return parameters cover the calibrated ones in this thesis. The suggested allocation in Iskhakov et al. (2015) is expected to be higher, owing to the constraint that all wealth is to be consumed. That aside, the result confirms the general findings in Iskhakov et al. (2015) and Bütler et al. (2016) – annuitisation is crowded out by the Age Pension and annuitisation increases with wealth but quickly flattens to a constant proportion. Both papers find evidence that the means-test impacts annuitisation, especially when binding. This can be seen as the decreasing annuitisation rate around $200,000 in Figure 6.2, which represents the transition from full to partial Age Pension. By annuitising at this (or lower) wealth level, no more Age Pension can be received by decreasing assets held, but the annuity payments lead to less Age Pension due to the income-test. When a partial pension is received, however, any annuity payments are only partly assessed in the income-test, hence annuitisation is still high until full Age Pension is received. The means-tested Age Pension thus effectively crowds out annuitisation at lower wealth, but not for wealthier households. There are no indications of high sensitivity to means-tested thresholds however, other than decreasing annuitisation rate when the means-test binds.

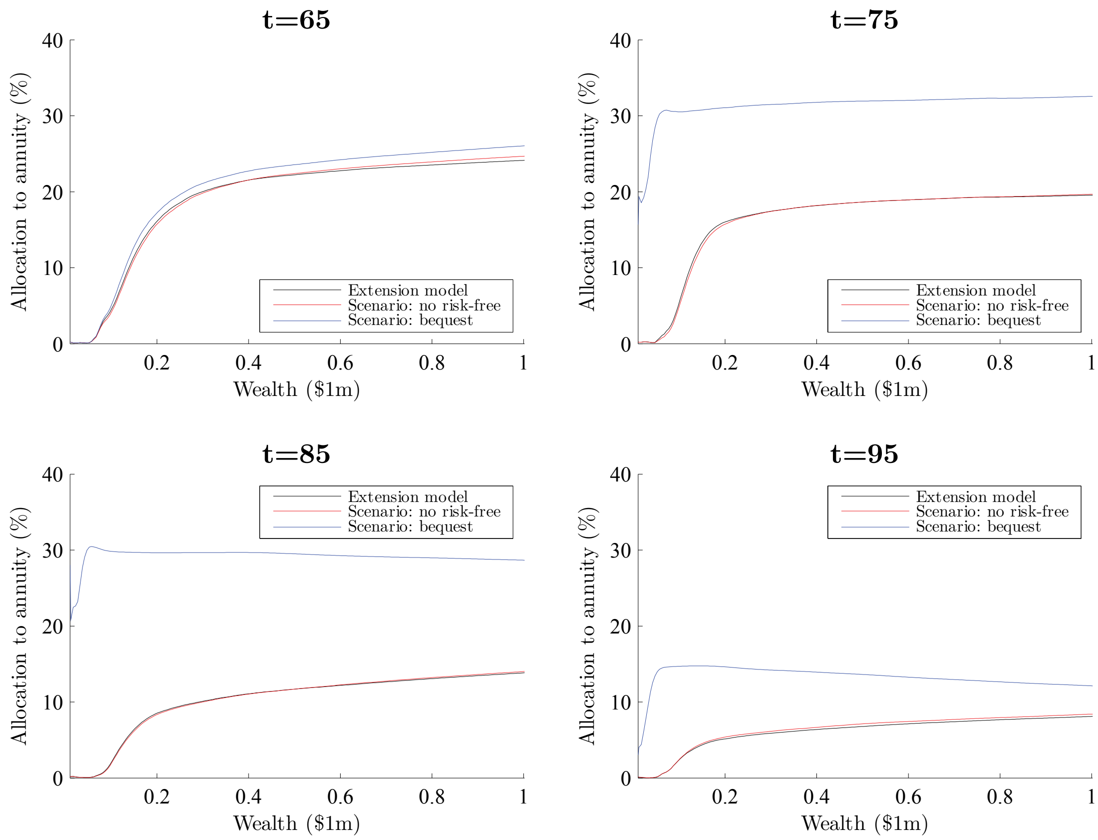

Figure 6.2: Optimal annuitisation at retirement given initial liquid wealth and no prior annuitisation. In addition to the standard case, the scenarios when no risk-free asset is available, and when the annuity value is included in the bequest, is shown.

One common argument against annuitisation is that wealth which could be bequeathed is given up. Therefore, the model was also solved where the current value of the annuity is added to the bequest function. The value was calculated from the current state of the annuity payment y_t and risk-free rate r_t using equation (6.19). The optimal annuitisation is shown by the dotted line in Figure 6.2. This scenario suggests that higher annuitisation is indeed optimal, where the difference is significantly higher a few years into retirement and especially for less wealthy households. The annuitisation rate is in fact higher for poor households in this case, as they benefit the most from additional utility from annuity payments and bequest, while the mortality credit keeps the annuity cost low. It should be noted that an annuity that has a residual value will have a higher price. This scenario was only to verify the significance of the argument that annuitisation is lower due to not being bequeathable.

Contrary to Iskhakov et al. (2015) and Bütler et al. (2016), but similar to Milevsky and Young (2007), the retiree is allowed to purchase annuities at any time, rather than only at time of retirement ($t = 65$). Figure 6.3 shows the total annuity allocation for a given initial liquid wealth during the retirement. In order to calculate this, it is assumed the retiree follows the optimal control and that wealth grows with the expected return. This gives a very different picture of optimal annuitisation than what is seen in Figure 6.2. Households with lower wealth now have a significant proportion of annuities. This is due to the effect of low consumption in older age, hence Age Pension payments accumulate and wealth increases, which is then partly annuitised. It is therefore sub-optimal for poor households to annuitise, but if their wealth grows it is indeed optimal to anuitise even at very late stages of retirement. The calculations of total annuitisation in retirement can also can be used to evaluate when in retirement annuitisation is optimal. The longer the retiree waits to annuitise, the larger the mortality credit will be in relation to price (due to the higher death probability), but on the other hand the desired excess consumption decreases towards the consumption floor. By deferring the choice to annuitise, the assets can instead be used to generate investment returns. Figure 6.4 shows the cumulative wealth allocated to annuities with age. The majority of total annuitisation only happens during the first five years of retirement, and then remains rather constant. This supports the findings in Milevsky and Young (2007) who shows it is optimal with immediate partial annuitisation, which also increases with wealth.

Early annuitisation indicates that it is not optimal to delay annuities in order to get increased risky exposure. There is one exception though - at very old age (around 95) it becomes optimal to add to the allocation. This is, however, due to growing wealth from decreased consumption with age, and not an effect of delaying annuitisation. Iskhakov et al. (2015) found that deferred annuities are more attractive to less wealthy retirees owing to the cheaper price. The extension model does not get the same result, due to the lack of additional mortality credit when using immediate annuities, compared to deferred annuities which are purchased before the annuity payments start.

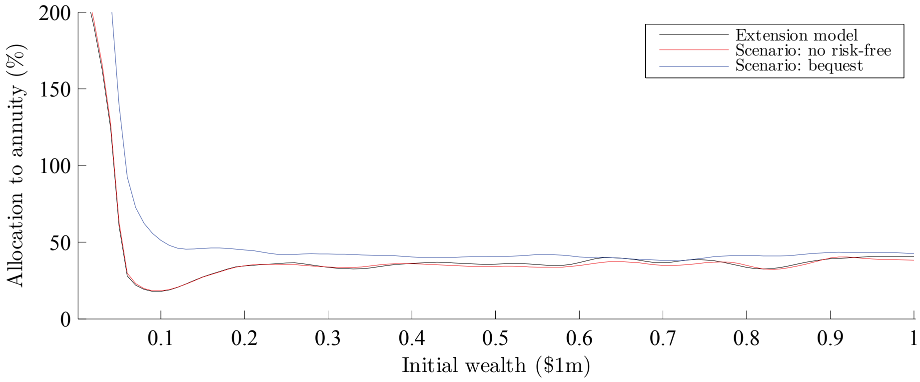

Figure 6.3: Optimal total allocation to annuities over the life time in retirement given initial liquid wealth.

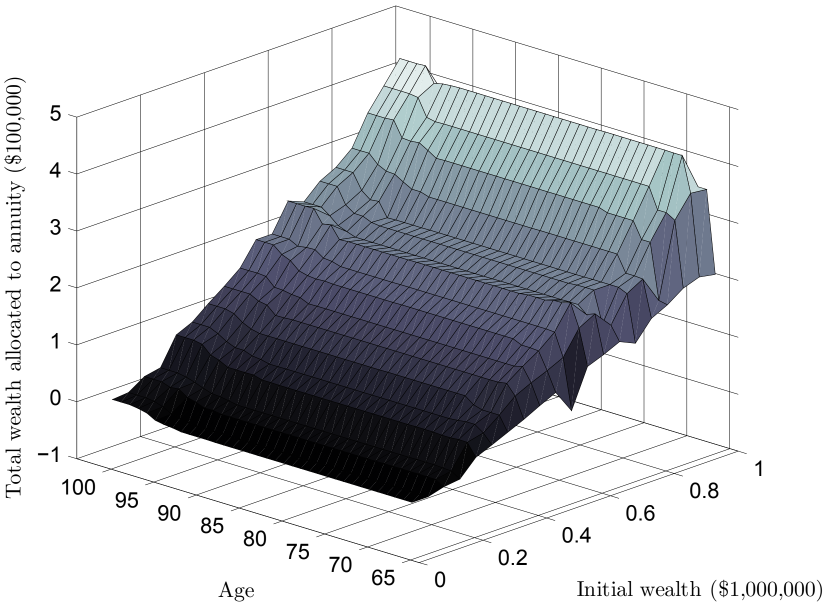

Figure 6.4: Optimal allocation to annuities over time in retirement given initial liquid wealth, assuming no previous annuitisation.

It should be noted that the optimal annuitisation is an upper bound due to the assumption of no commercial loadings. If a commission or management fee was present, this would make the annuity even less desirable. In addition, since wealthier households tend to live longer than less wealthy (De Nardi et al., 2010), the annuitisation is potentially underestimated for the wealthier households and overestimated for the less wealthy household. As the model does not include medical expenses at older age, nor aged care, it can be argued that additional annuitisation is optimal when these costs are included. At the same time, since entering aged care (i.e. a retirement village) attracts rather large one-time costs, this can decrease the optimal level of annuitisation. The finding that annuitisation is optimal only early in retirement might also change in this case.

6.4.2 Extension 2: Scaling housing

The purpose of this extension of the model is to evaluate whether scaling housing, or accessing home equity, is optimal in retirement. In order to test this, it is important that the retiree starts with the optimal house asset at the time of retirement. If not, then the solution might suggest scaling housing just to meet the initial optimal ratio of house assets to liquid wealth. This does not reflect whether it is optimal to scale housing in retirement however, only that it is optimal with a certain level of housing assets in relation to wealth once retired. The retiree therefore starts with the optimal house asset at retirement for a given liquid wealth, and the wealth paths and optimal control is then simulated until terminal time T. Figure 6.5 shows the wealth, housing and reverse mortgage paths throughout retirement based on optimal decisions and the expected return on risky assets. Three different levels of total initial wealth at retirement are considered: \$1m, \$2m and \$4m where it is optimal to allocate approximately 80%, 77.5% and 75% respectively into housing for a single household. The single household is chosen as the relative risk aversion for housing is slightly lower than risk aversion for consumption. As can be seen, in none of the cases is it optimal to downscale housing, while all of them take advantage of the reverse mortgage to keep liquid wealth at a relatively constant level. The loan value is added on to during retirement when required, but also grows based on the interest accumulated.

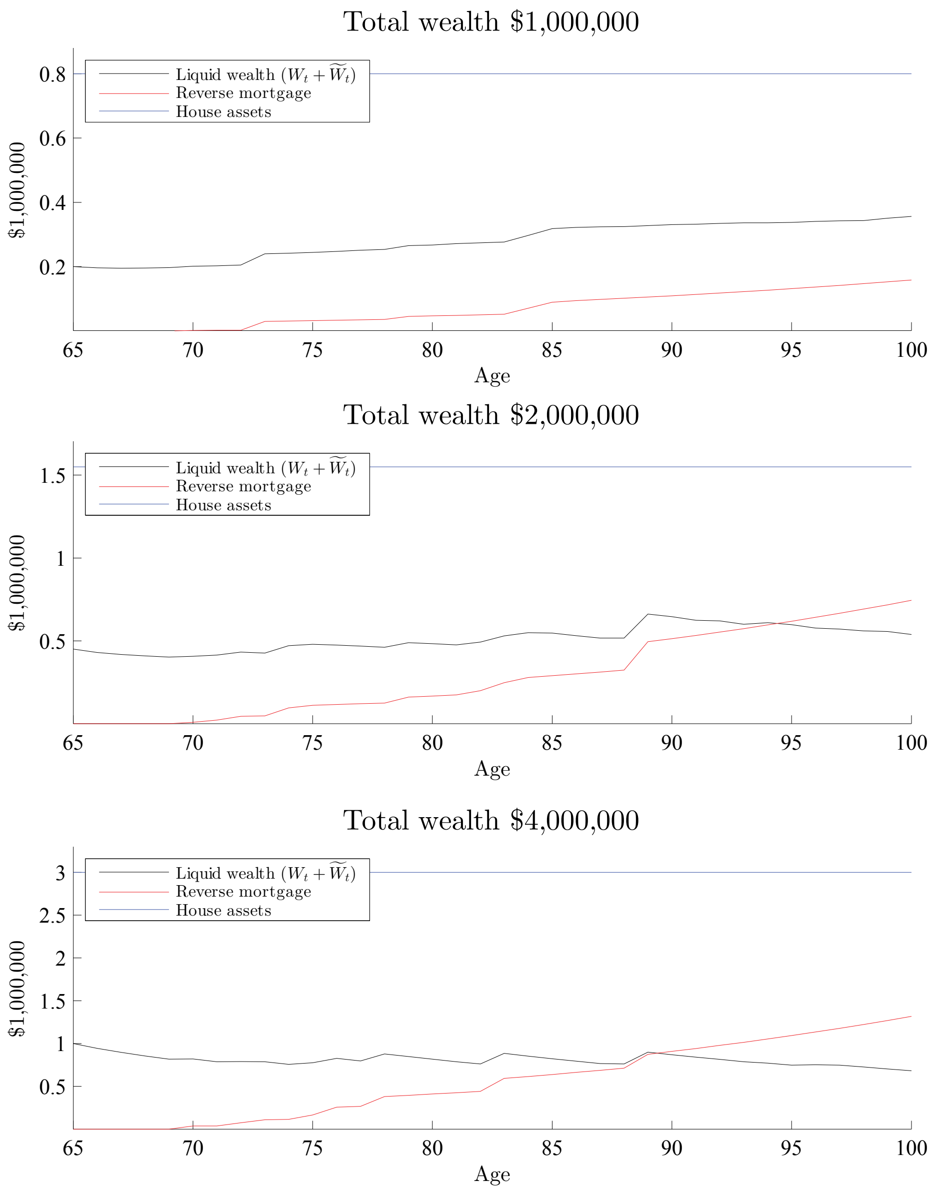

Figure 6.5: Wealth, house and reverse mortgage paths in retirement given low, medium and high initial total wealth.

The optimal reverse mortgage as a proportion of the house value decreases with wealth, and increases with the house value. Irrespective of house value, the loan proportion starts at the same value for households with no wealth. One might expect that the proportion would be less for a higher house value, as this would still access more wealth for the retiree, but this is not the case. However, the higher the house value, the more liquid wealth the retiree can have and still optimally take out a reverse mortgage. This confirms the results in Chiang and Tsai (2016), who found that as age increases, and the higher the initial wealth and house price is, the more willing the retiree is to use reverse mortgages. Figure 6.6 shows the optimal loan proportion for different house values in relation to liquid wealth for single households, where the proportion in relation to wealth has a very linear relationship. A less wealthy household, which might need the wealth more than a wealthier household, generally should not take out a reverse mortgage unless the house value is substantially higher than the liquid wealth. Each line in Figure 6.6 reaches zero before it equals the optimal liquid wealth given the house value, hence a reverse mortgage is never optimal until wealth is drawn down enough to differ significantly from the house asset. The same relationship holds true for couple households, although at a slightly higher wealth level than for singles. When comparing the optimal loan proportion over time in retirement, the initial maximum level of approximately 10% increases yearly, but flattens out around year 80 and then remains constant at approximately 20%. The LVR threshold therefore never binds when a loan is created, given the calibrated parameters. It is reasonable to expect that if the risk aversion or preferences for bequest decreases, then the optimal loan value might increase. The optimal reverse mortgage in the solution is also an upper bound, as additional commercial loadings such as a fee to initiate the loan might apply in reality. However, a reverse mortgage could theoretically be refinanced if interest rates drop, thus any costs associated with the loan can be lowered that way.

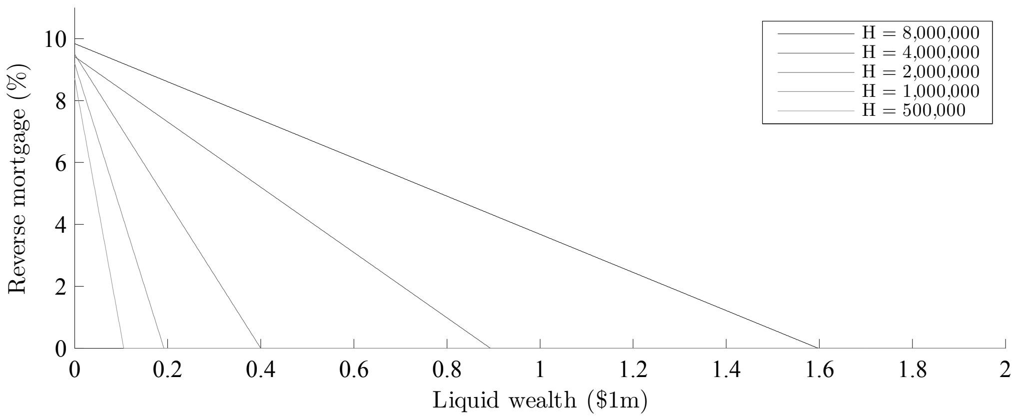

Figure 6.6: Optimal proportion of reverse mortgage given housing wealth and liquid wealth at retirement for a single household.

If the retiree's housing asset is significantly less than optimal, then the solution will quickly suggest that the retiree should scale (or acquire) housing assets to get close to the optimal level. However, the opposite does not hold true. If the retiree starts with housing assets significantly larger than optimal, then it is not optimal to downscale, with the exception if the retiree has close to no wealth at all but significant wealth in the house asset. In general, it is therefore never optimal to downscale housing in retirement[10], not even when reverse mortgages are not available. Only in the case of an event which would incur a significant cost, such as a medical issue, would it be reasonable to downscale. This event is not modelled, however, and would be a result of the budget constraints due to the threshold of the loan value, rather than to maximise utility. It is not optimal to upscale housing once retired either, with the exception of very low house assets (∼$100,000 or less) which only reflects the desire to get close to the initial optimal ratio rather than an actual upsizing decision. The reason why downsizing housing is not optimal stems from a combination of the high cost associated with the sale of the house while housing is included in the bequest (hence wealth is given up by downsizing), and that the calibrated consumption floor is already covered by Age Pension payments. If the retiree wants to access just part of his/her house wealth, then downsizing the house will first incur a transaction cost on the full home value, even if the retiree only downscales slightly. To access 10% of the housing wealth, he/she needs to give up 6% of this equity in costs. It is therefore much more economical to take out a reverse mortgage. At the same time, housing utility is received based on the value of the house, even if there is an outstanding reverse mortgage. By utilising the reverse mortgage the retiree can therefore keep a high housing utility, while still accessing the equity. The retiree will give up bequeathable wealth as the loan value accumulates interest, but the funds received from the reverse mortgage can either be invested at a higher (although risky) return and the loss of utility in bequest is partly compensated with higher housing utility through retirement.

A final interesting outcome is that the additional decision variables impact the optimal housing in relation to wealth. With access to the housing asset in retirement, it is now optimal to allocate slightly more towards housing as can be seen in Figure 6.7. This effect is due to the possibility of 'hiding' away assets in a family home, and then tapping into these assets using the reverse mortgage. It should also be noted that unlike other jurisdictions, Australia does not tax the imputed rent of

[10] It should be noted that we only consider the case with no outstanding mortgage on the house. If a significant mortgage exist, the outcome may differ.

housing, further adding to the bias towards holding housing as an asset. A retiree can avoid having assets included in the means-test by over allocating to housing assets, and therefore receive additional 'free' wealth from the Age Pension. As the liquid wealth is consumed, it can be replenished by taking out a reverse mortgage, while still accessing the Age Pension.

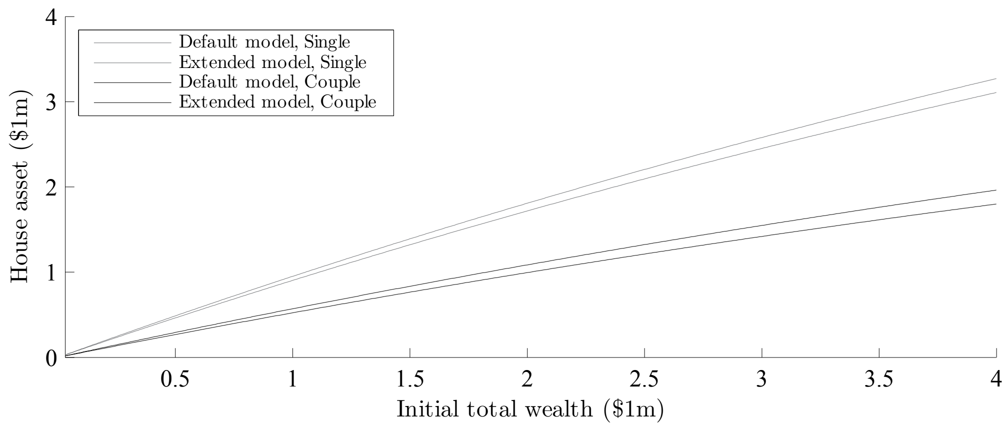

Figure 6.7: Optimal allocation to housing at retirement for the default case compared to extension model 2 where decisions for scaling housing and reverse mortgage are available.

6.5 Conclusion

In this chapter the retirement model was extended with the option to annuitise wealth, and to scale housing or access the home equity with a reverse mortgage. It was then evaluated whether such options were optimal during retirement in relation to the means-tested Age Pension.

In general, the optimal annuitisation in a realistic retirement model verifies previous research performed with more restricted models. The means-tested Age Pension crowds out annuitisation, and the alternative to allocating wealth to an annuity is preferred over the risk-free rate. Even when a partial Age Pension is received it is optimal with partial annuitisation, although decreases quickly around the threshold for the full Age Pension. For wealthier households, the annuity payments are much higher than the partial Age Pension received, so even if 'free' wealth is given up the retiree is better off annuitising. The total allocation to annuities, or the point in time, is not known at the time of retirement and depends on the outcome of stochastic factors in retirement.

An annuity provides a significant discount in terms of mortality credits, where the additional utility is higher compared with the alternative to invest the funds in risky assets and annuitise at a later stage in retirement. As consumption decreases with age, this could make annuitisation less desirable, but the results do not indicate this to be true as the mortality credit is higher at older age. It is optimal to annuitise sooner rather than later as it is cheaper to store wealth in an annuity rather than risk-free investments. In the Australian setting, it is not optimal with a one-off decision to annuitise, but rather to gradually increase allocation in the first five years in retirement, and to annuitise additional wealth later in retirement depending on the wealth evolution.

A retiree is in general better off utilising a reverse mortgage rather than downsizing the house, despite the accumulated interest of the loan. By keeping a house that is larger than optimal while drawing down the housing assets, the retiree still receives utility from living in the house, while it is still partly bequeathable. The additional utility from this outweighs the cost of an outstanding reverse mortgage. A reverse mortgage does therefore not necessarily benefit a retiree financially, unless the retiree can access additional Age Pension payments by 'hiding' assets in the family home, but it does help maximising utility throughout retirement. The optimal decisions are, however, subject to wealth levels and housing assets, where wealthier retirees with more housing assets optimally access a higher proportion reverse mortgage than less wealthier households.

Chapter 7

Conclusion

In this thesis we model the retirement phase of the life cycle for Australian retirees. The model presented is more realistic than currently available models with respect to stochastic factors and control variables. It allows for sequential family status, home equity in the bequest function, and a "health" proxy to represent declining consumption in retirement. The model is shown to capture the characteristics of Australian empirical data well, and is further fitted to the data via a maximum likelihood calibration to the 'Household Expenditure Survey' and the 'Survey of Income and Household' from Australian Bureau of Statistics (2011). This allows us to evaluate the optimal behaviour for an individual retiree, given a set of circumstances (family status, wealth level, age, homeowner status) in relation to the Australian Age Pension. We evaluate such behaviour under four recent policy changes.

A limitation of the model is the fact that it requires a numerical solution, which stems for the piecewise linear Age Pension function. Therefore, in order to extend the model with additional states and control variables we utilise the Least-Squares Monte Carlo method, which is an approximate dynamic programming method. This method involves a forward simulation and backward optimisation, and is known to have drawbacks when the time period of the problem increases and when the objective function is characterised by utility functions. To overcome the difficulties with longer time periods we adopt the approach in Kharroubi et al. (2014, 2015) and improve the exploration of the state. To handle the bad fit of the utility based objective function, a transformation is proposed of the value function where a bias function is used to manage the re-transformation bias otherwise present. Such a transformation when accounting for the bias is also shown to handle the case when optimal control affects the disturbance term, which otherwise will not be reflected properly and risk then becomes underestimated.

The Least-Squares Monte Carlo approach is shown to work well, and allows the model to be extended with annuitisation and access to home equity through up/downscaling of the home or by taking out a reverse mortgage. The analysis of the results with respect to the model sheds new light on previous research, that was in more restrictive frameworks. It is optimal with partial annuitisation in retirement, even when the means-test binds, although at a lower rate. The decision to annuitise should be taken early in retirement, as deferring the decision leads to a utility loss even if the mortality credit is higher. With respect to housing decisions, it is generally never optimal to scale housing in retirement, other than to meet the initial optimal house to wealth ratio early in retirement. Downscaling would only occur due to the budget constraint, where the retiree has no other option in order to cover costs. A reverse mortage, however, is very beneficial for an Australian retiree, as it allows the retiree to receive utility from the home while still accessing home equity and partly bequeathing the home. Although such a loan can be expensive, the model suggests that all retirees are better off utilising one as the wealth-to-house ratio decreases due to wealth drawdown.

7.1 Major findings

The major findings in the research project were:

- Demonstration that the sequential model developed can capture the main characteristics of the typical Australian household in retirement. A time-dependent declining preference for consumption, modelled as a "health" proxy, captures the shape of empirical consumption for given wealth and age.

- Demonstration that the model is well calibrated to Australian empirical data, measured in relation to consumption (given age) and to housing assets.

- Finding that an Australian retiree is better off with the recent Age Pension policy changes. Despite steeper taper rates, the average retiree will receive more Age Pension through retirement owing to the move from assessing drawn down wealth as income to deemed income. In addition, a retiree can optimally allocate even more to housing under the active Age Pension policy.

- Demonstration that Least-Squares Monte Carlo can be applied to utility function based life cycle modelling. The accuracy is greatly increased by regressing on the realised value function rather than the regression surface, which minimises the risk of regression errors growing large when the number of time periods in the problem increases.

- Development of an approach for bias-correction in order to decrease the regression errors when using Least-Squares Monte Carlo. The common transformation of the utility function, which improves the linearity of the value function to help with the regression, underestimates any risk (stochasticity) in the solution. In addition, control variables related to stochastic variables are therefore incorrect. To cope with this it was shown that a re-transformation bias function can correct both the underestimated risk and the optimal control of such variables.

- Finding that it is indeed optimal for Australian retirees to partly annuitise wealth, even when the means-test is binding, although at a lower rate. The Age Pension does crowd out annuitisation, but not completely.

- Finding that it is optimal to over allocate assets to housing in retirement, and instead to access the home equity through a reverse mortgage. It is generally not optimal for the Australian retiree to scale house assets.

7.2 Applications

The model presented in the thesis has several areas of applications:

- To evaluate new financial products in retirement. The model can be extended or adapted to evaluate how retirees' behavior can change in relation to such products, or whether it changes the welfare measured in terms of utility or lifetime wealth. For example, we found that partial annuitisation is optimal for most retirees, and that reverse mortgages are preferred over downsizing. Based on this, a financial product which pays an income stream and is paid for (partially) with the home might be of interest. Such a product could also have a health insurance component if the retiree must move to an aged care facility.

- To evaluate the impact of policy changes. It is very difficult for the government to estimate budget needs and welfare changes for proposed policy changes. The effect at a micro level is not known until empirical data exist after implementation. A policy change that is expected to distort the behaviour in a particular way might not have the expected outcome if the retiree has sufficient freedom in his/her decisions. The framework in this thesis can therefore improve the understanding for policy-makers of proposed policy changes. This can be linked to the next bullet point.

- To model effects on a macro level by aggregating data from micro-simulations. The model can be solved with the solution spread over a wide range of state values. The (retired) population in general can be represented by a large number of samples, distributed in the same way as the population distribution in respect to the state variables. Each sample is then simulated by associating it with a starting point in the model solution, and the evolution of state variables is simulated by using the optimal control from the solution for each time period. These aggregated simulations can help in understanding the implications of, e.g., policy changes on a macro scale.

- From an international perspective, there are many application areas as well. The calibrated values are expected to be similar in other developed countries owing to the similar characteristics; hence, their specific pension systems can be modelled by adjusting the Age Pension function and necessary constraints. The findings presented here can spark ideas for international policy-makers as well. For example, the fact that it is optimal for retirees to accept additional risk could be beneficial if a country wants to encourage investments, as pension savings and equity investments would provide a significant inflow of capital to the country's markets.

- The proposed method to improve the accuracy of utility functions used in Least-Squares Monte Carlo extends to more areas, such as portfolio optimisation, game theory or behavioural economics and general decision-making. It is also applicable to prospect theory problems.

7.3 Further study

The model presented has many benefits, but also a number of limitations which can open the way to further study. For example, the data used in the calibration does not include samples in assisted care facilities, hence this 'phase' in retirement was not explicitly modelled. It is likely that retirees living in these facilities face significant costs for living arrangement as well as medical costs. A suitable extension to the model would be to add another state to the family status, and to account for these expenses.

Related to the data limitation above, the calibration does not allow for cohort effects. To improve the calibration further it would therefore be of interest to calibrate the model on waves of data that allow for cohort changes to be included. This allows for a specifically interesting analysis, which is how retirement behaviour changes with respect to a financial crises, as recent data will cover a couple of them. However, calibration over time poses additional difficulties as policy changes have been rather frequent recently. This leads to a time dependent Age Pension function, which applies to different age groups to reflect the evolution of the Age Pension policies.

Further options and parameterisation can be introduced in the model, based on the findings in the thesis. For example, we found that it is not optimal to downscale housing in retirement. As the utility from owning a house is constant in retirement, but the retiree will not require the same size of house at older age due to reasons such as living alone or difficulties moving around, it might be reasonable to assume a decreasing housing utility with age.

Finally, the model can be extended with the accumulation phase leading up to retirement. Such an extension would include variables related to salary and education which affect the total accumulation of Superannuation assets, and can allow for voluntary retirement decisions. A lot of research exists on this front already, but it is limited with regards to Superannuation rules which affect the constraint and outlook for asset accumulation during the life cycle.

Appendix A

Data aggregation

The data sample is taken from the 'Household Expenditure Survey' and the 'Survey of Income and Household' 2009-2010 (Australian Bureau of Statistics, 2011). The information is available on both Person level and Household level.

Data are taken from the Personal level, where the Personal level is matched to the household data and expenditure data via the sample IDs. Up to two Personal level samples can therefore be matched to one household sample. The data are then filtered on Age $\geq$ 27 (class 27 indicates 65 years old) (AGERHBC) and Labour status (3 = Not in Labour force) (LFSRH). Single/Couple classification is made from Income Unit Type (IUTYPEP), where 1-2 represents a Couple and 3-4 Single household.

Since the data are spread between Person level, Household level, and Expenditure level, the different levels need to be matched together by the associated ID. For couples, the sample set is created from merging the individual (Person level) assets and liabilities, and then adding this to the Household level data. For singles, the Household level data and Person level data can be matched one-to-one. Expenditure level data is summed up based on Household ID to match with Household level data, so that the final sample data have all variables on Household level. The data are then aggregated to reflect the variables of interest for the calibration:

Consumption Variable aggregated from Expenditure level variables.

All categories are used, with the exception of:

Category 17, 'Superannuation and Life Insurance'.[1]

Category 15, 'Mortgage repayments'.[2]

Pension Variable aggregated from Personal level pension variables:

'Current weekly income from age pension' (IAGECP)

'Current weekly income from service pension (DVA)' (ISERVCP)

'Current weekly income from disability support pension' (IDSUPPCP)

'Current weekly income from pension supplements' (IPSUPCP)

'Current weekly income from disability pension' (IDISBCP)

Family home Variable aggregated from Household level:

'Estimated sale price of dwelling' (HVALUECH)

Minus 'Principal outstanding on selected dwelling' (LIASDCH).

Wealth Variable aggregated from Person level wealth variables:

'Balance of accounts with government superannuation funds' (VSUPGCP)

'Balance of accounts with non-government superannuation funds' (VSUPNCP)

'Value of accounts held with financial institutions' (VFINCP)

'Value of debentures and bonds' (VDEBCP)

'Value of loans to persons not in the same household' (VPLNCP)

'Value of other financial investments' (VINVOTCP)

'Value of shares - person level' (VSHARCP)

'Value of public unit trusts - person level' (VPUTTCP)

'Value of own incorporated business (net of liabilities)' (VIBUSCP)

'Value of own unincorporated business (net of liabilities)' (VUBUSCP)

'Value of silent partnerships - person level' (VSIPCP)

'Value of private trusts - person level' (VPRTCP)

Plus Household level wealth:

'Net wealth of household' (WEALTHH)

[1]The reason is that income that is not consumed during retirement will automatically be saved, hence counting this as an expense would be incorrect.

[2]This variable accounts only for the principal component, where the effect of mortgage repayments are already reflected in the assets/liabilities.

‘Value of residential property excl selected dwelling’ (VRPRCH)

‘Value of non-residential property’ (VNRPRCH)

Minus Household level liabilities, children’s wealth, and vehicles/housing contents:

‘Principal outstanding on rental property loans’ (LIARPCH)

‘Principal outstanding on loans for other property’ (excl business and investment loans) (LIAOPCH)

‘Amount of credit card debt - household level’ (LIACCCH)

‘Amount of HECS/HELP liability’ (LIAHECCH)

‘Amount of Student Financial Supplement liability’ (LIASFSCH)

‘Principal outstanding on investment loans’ (excl business and rental property loans) (LIAINVCH)

‘Principal outstanding on loans for vehicle purchases (excl business and investment loans’ (LIAVECH)

‘Principal outstanding on loans for other purposes’ (excl business and investment loans) (LIAOTCH)

‘Principal outstanding on selected dwelling’ (LIASDCH)

‘Value of children’s assets’ (VCHASSCH)

‘Value of vehicles’ (VVEHICH)

‘Value of contents of selected dwelling’ (VCONTCH)

To clean the data from outliers and incorrectly reported variables, the following final steps are taken:

- Remove samples with yearly expenditure less than $3000.

- Remove samples where yearly expenditure is larger than half the wealth and three times the Age Pension received.

- Remove samples where no Age Pension is received, although the means-test indicates significant pension payments should be received.

- Remove samples where both pension and consumption were zero.

For singles this left 2,038 samples, and for couples 2,017 samples.

Appendix B

Duan's Smearing Estimate

The results in this section are based on Duan (1983). Denote the non-transformed observations $Y_i, i = 1, ..., n$ and the transformed observations $\eta_i, i = 1, ..., n$ such that $\eta_i = g(Y_i)$, $Y_i = h(\eta_i)$, i.e. $h := g^{-1}$. Assume g (the transformation) and h (the re-transformation) are known monotonic and continuously differentiable functions, such as a CRRA utility function $h(x) = x^\gamma/\gamma,\ \gamma < 0$. Consider the linear regression carried out on the transformed observations

$$\eta_i = \boldsymbol{\beta}'\boldsymbol{X}_i + \epsilon_i, \quad \epsilon_i \overset{i.i.d}{\sim} F(\cdot), \quad \mathbb{E}[\epsilon_i] = 0, \quad \mathrm{var}[\epsilon_i] = \sigma^2, \tag{B.1}$$

where $\boldsymbol{\beta}$ is the vector of coefficients, $\boldsymbol{X}_i$ is the vector of covariates and ϵ_i are the independent and identically distributed residuals from some zero mean distribution $F(\cdot)$ with finite variance. The error terms do not need to have a known distribution, although they are expected to have zero mean and constant variance. Now, if the re-transformation is applied to the prediction of the transformed variables we would get an incorrect estimate due to Jensen's inequality, because $\mathbb{E}[Y] \leq h(\mathbb{E}[\boldsymbol{\beta}'\boldsymbol{X} + \epsilon])$ if h is a concave function such as a utility function.

Smearing Estimate attempts to approximate the non-transformed expectation

$$\mathbb{E}[Y] = \mathbb{E}\left[h(\boldsymbol{\beta}'\boldsymbol{X} + \epsilon)\right] = \int h\left(\boldsymbol{\beta}'\boldsymbol{X} + \epsilon\right) dF(\epsilon) \tag{B.2}$$

after estimating the regression coefficients $\widehat{\boldsymbol{\beta}}$ using the empirical distribution function of the residuals $\widehat{\epsilon}_i = \eta_i - \widehat{\boldsymbol{\beta}}'\boldsymbol{X}_i$:

$$\widehat{F}_n(e) = \frac{1}{n}\sum_{i=1}^{n} \mathbb{I}\{\widehat{\epsilon}_i \leq e\}, \tag{B.3}$$

where $\mathbb{I}\{\cdot\}$ is the indicator symbol that equals 1 if the statement in brackets $\{\cdot\}$ is true and 0 otherwise. The estimated expectation of Y can then be found as

$$\widehat{\mathbb{E}}[Y] = \int h\left(\widehat{\boldsymbol{\beta}}'\boldsymbol{X} + \epsilon\right) d\widehat{F}_n(\epsilon) = \frac{1}{n}\sum_{i=1}^{n} h\left(\widehat{\boldsymbol{\beta}}'\boldsymbol{X} + \widehat{\epsilon}_i\right). \tag{B.4}$$

To illustrate, suppose we consider regression $\ln Y_i = \boldsymbol{\beta}'\boldsymbol{X}_i + \epsilon_i$ and we want to estimate $\mathbb{E}[Y^\gamma/\gamma]$, then the Smearing Estimate is

$$\frac{1}{n}\sum_{i=1}^{n} \frac{\left(e^{\widehat{\boldsymbol{\beta}}'\boldsymbol{X}+\widehat{\epsilon}_i}\right)^\gamma}{\gamma} = \frac{\left(e^{\widehat{\boldsymbol{\beta}}'\boldsymbol{X}}\right)^\gamma}{n\gamma}\sum_{i=1}^{n} e^{\widehat{\epsilon}_i\gamma}. \tag{B.5}$$

The Smearing Estimate works well for non-normal errors and can accommodate for heteroskedasticity, provided it is not related to a covariate (Duan, 1983).

Appendix C

Controlled Heteroskedasticity

Consider a simple model with heteroskedasticity, such as $Y = \boldsymbol{\beta}'\boldsymbol{X} + \epsilon$ where $\boldsymbol{X}$ is a vector of covariates, $\boldsymbol{\beta}$ is a vector of regression coefficients, $\mathbb{E}[\epsilon] = 0$ and $\text{var}[\epsilon] = \sigma^2 c(\boldsymbol{X})$. There are various ways to estimate function $c(\boldsymbol{X})$ that is causing heteroskedasticity. In particular, we adopt a popular method from Harvey (1976) (also see Baser (2007), (Greene, 2008, chapter 8)). Assume $c(\boldsymbol{X}) = e^{\boldsymbol{\mathcal{L}}'\boldsymbol{X}}$ to avoid negative values, where $\boldsymbol{\mathcal{L}} = \mathcal{L}_0, \mathcal{L}_1, ..., \mathcal{L}_K$ is another vector of regression coefficients. Thus

$$\epsilon^2 = \sigma^2 c(\boldsymbol{X}) v = \sigma^2 e^{\boldsymbol{\mathcal{L}}'\boldsymbol{X}} v^2, \quad \mathbb{E}[v] = 0,\ \mathbb{E}[v^2] = 1 \tag{C.1}$$

and we can write

$$\ln(\epsilon^2) = a + \mathcal{L}_1 X_1 + ... + \mathcal{L}_K X_K + \ln v^2, \tag{C.2}$$

where $a = \ln(\sigma^2) + \mathcal{L}_0$. The parameter estimates are found by two-stage procedure. First, we find the ordinary least squares estimate $\widehat{\boldsymbol{\beta}}$ and calculate the observed residuals $\widehat{\epsilon} = Y - \widehat{\boldsymbol{\beta}}'\boldsymbol{X}$. Then we perform the ordinary linear regression (C.2) where unobserved ϵ are replaced with $\widehat{\epsilon}$ to estimate the variance function $\widehat{\sigma}^2 \exp(\widehat{\boldsymbol{\mathcal{L}}}'\boldsymbol{X})$. Finally, using estimated variance, $\boldsymbol{\beta}$ is approximated by the weighted least squares method. The process can be iterated to improve the estimates.

Other methods to estimate $c(X)$ include random effect representation (Hoff and Niu, 2012), kernel estimates (Muller and Stadtmuller, 1987) or via link functions (Smyth, 1989).

Appendix D

Solution to multi-period utility model

In this section we derive the analytical solution for optimal drawdown and risky asset allocation in the multiperiod utility model considered in Section 5.5.2. The objective is to maximize the expected value function

$$V_0(x) = \sup_{\pi} \mathbb{E}\left[\sum_{t=0}^{N-1} \frac{(\alpha_t X_t)^\gamma}{\gamma} \;\Big|\; X_0 = x; \pi\right]. \tag{D.1}$$

Let ξ_t represent the stochastic component in the transition function, such as $\xi_t = e^{Z_{t+1}}$ in the case of a single risky asset. This type of problem was originally solved in Samuelson (1969).

At the terminal time $t = N$, the value function is given by

$$V_N(X_N) = \frac{(\alpha_N X_N)^\gamma}{\gamma}. \tag{D.2}$$

It is optimal to consume all wealth as no utility is received from saving wealth, hence by intuition $\alpha_N = 1$. The risky asset allocation at this point has no impact.

At time $t = N - 1$, the value function is

$$\begin{aligned} V_{N-1}(X_{N-1}) &= \frac{(\alpha_{N-1} X_{N-1})^\gamma}{\gamma} + \mathbb{E}\left[V_N(X_N)\right] \\ &= \frac{(\alpha_{N-1} X_{N-1})^\gamma}{\gamma} + \frac{((1-\alpha_{N-1})X_{N-1})^\gamma \mathbb{E}\left[\xi_{N-1}^\gamma\right]}{\gamma}. \end{aligned} \tag{D.3}$$

To find the optimal drawdown, differentiate with respect to α_{N-1}

$$\frac{\partial V_{N-1}}{\partial \alpha_{N-1}} = X_{N-1}(\alpha_{N-1}X_{N-1})^{\gamma-1} - X_{N-1}((1-\alpha_{N-1})X_{N-1})^{\gamma-1}\mathbb{E}\left[\xi_{N-1}^{\gamma}\right], \tag{D.4}$$

set this equal to 0 and solve for α_{N-1}

$$\begin{aligned}
&X_{N-1}(\alpha_{N-1}X_{N-1})^{\gamma-1} - X_{N-1}((1-\alpha_{N-1})X_{N-1})^{\gamma-1}\mathbb{E}\left[\xi_{N-1}^{\gamma}\right] = 0\\
&\Rightarrow \alpha_{N-1} = (1-\alpha_{N-1})\mathbb{E}\left[\xi_{N-1}^{\gamma}\right]^{\frac{1}{\gamma-1}}\\
&\Rightarrow \alpha_{N-1} = (1+\mathbb{E}\left[\xi_{N-1}^{\gamma}\right]^{\frac{1}{1-\gamma}})^{-1}.
\end{aligned} \tag{D.5}$$

If the stochastic growth of wealth depends on a control variable, such as if $\xi_{N-1} = e^{\delta_{N-1}Z_N+(1-\delta_{N-1})r_{N-1}}$ considered in the model in section 5.5.2, then the same steps are used to find the optimal risky asset allocation δ_{N-1}. In this case, assuming $Z_t \overset{i.i.d}{\sim} \mathcal{N}(\mu, \sigma^2)$,

$$\frac{\partial V_{N-1}}{\partial \delta_{N-1}} = \mathbb{E}\left[(Z_N - r)((1-\alpha_{N-1})X_{N-1})^{\gamma}\xi_{N-1}^{\gamma}\right] \tag{D.6}$$

$$\begin{aligned}
&\Rightarrow \mathbb{E}\left[(Z_N - r)(e^{\delta_{N-1}Z_N+(1-\delta_{N-1})r_{N-1}})^{\gamma}\right] = 0\\
&\Rightarrow \delta_{N-1} = \frac{r-\mu}{\gamma\sigma^2}.
\end{aligned} \tag{D.7}$$

Finally, use α_{N-1} to find the maximum of the value function V_{N-1}

$$\begin{aligned}
V_{N-1}(X_{N-1}) &= \frac{(\alpha_{N-1}X_{N-1})^{\gamma}}{\gamma} + \frac{((1-\alpha_{N-1})X_{N-1})^{\gamma}\mathbb{E}\left[\xi_{N-1}^{\gamma}\right]}{\gamma}\\
&= \frac{(X_{N-1})^{\gamma}}{\gamma}((\alpha_{N-1})^{\gamma} + (1-\alpha_{N-1})^{\gamma})\mathbb{E}\left[\xi_{N-1}^{\gamma}\right]\\
&= \frac{(X_{N-1})^{\gamma}}{\gamma}(\alpha_{N-1})^{\gamma-1},
\end{aligned} \tag{D.8}$$

which will be used in the next iteration. By repeating these steps for $t = N-2, ..., 0$ a distinct pattern is found, where

$$\alpha_t = \begin{cases} 1, & \text{if } t = N,\\ (1+(\mathbb{E}\left[\xi_t^{\gamma}\right]\alpha_{t+1}^{\gamma-1})^{\frac{1}{1-\gamma}})^{-1}, & \text{otherwise.}\end{cases} \tag{D.9}$$

Bibliography

Agnew, J., H. Bateman, and S. Thorp (2013). Superannuation Knowledge and plan behaviour. *The Finsia Journal of Applied Finance* (1), 45–50.

Aïd, R., L. Campi, N. Langrené, and H. Pham (2014). A Probabilistic Numerical Method for Optimal Multiple Switching Problems in High Dimension. *SIAM Journal on Financial Mathematics 5*(1), 191–231.

Ameriks, J., A. Caplin, S. Laufer, and S. Van Nieuwerburgh (2011). The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Public Care Aversion from Bequest Motives. *The Journal of Finance 66*(2), 519–561.

Andréasson, J. G. and P. V. Shevchenko (2017). Assessment of Policy Changes to Means-Tested Age Pension Using the Expected Utility Model: Implication for Decisions in Retirement. *Risks 5*(3).

Andreasson, J. G. and P. V. Shevchenko (2017a). Bias-corrected Least-Squares Monte Carlo for utility based optimal stochastic control problems. *Preprint SSRN: 2985828*. Available at https://ssrn.com.

Andreasson, J. G. and P. V. Shevchenko (2017b). Optimal annuitisation, housing decisions and means-tested public pension in retirement. *Preprint SSRN: 2985830*. Available at https://ssrn.com.

Andreasson, J. G., P. V. Shevchenko, and A. Novikov (2017). Optimal Consumption, Investment and Housing with Means-tested Public Pension in Retirement. *Insurance: Mathematics and Economics 75*, 32–47.

ASFA (2015). Superannuation Statistics. URL http://www.superannuation.asn.au/ArticleDocuments/129/SuperStats-August2015.pdf.aspx (Accessed 2015-10-03).

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

Australian Bureau of Statistics (2011). Household Expenditure Survey and Survey of Income and Housing Curf Data. URL http://www.abs.gov.au/ausstats/abs@.nsf/mf/6503.0. (Accessed 2014-06-06).

Australian Bureau of Statistics (2012). Life Tables, States, Territories and Australia, 2009-2011. URL http://www.abs.gov.au/AUSSTATS/abs@.nsf/DetailsPage/3302.0.55.0012009-2011 (Accessed 2016-11-04).

Australian Bureau of Statistics (2014). 3101.0 - Australian Demographic Statistics, June 2014. URL http://www.abs.gov.au/ausstats/abs@.nsf/mf/3101.0 (Accessed 2016-11-14).

Australian Bureau of Statistics (2015). 3310.0 - Marriages and Divorces, Australia, 2014. URL http://www.abs.gov.au/ausstats/abs@.nsf/mf/3306.0.55.001 (Accessed 2016-11-14).

Australian Department of Treasury (2010). Australia to 2050: future challenges. Technical report, Department of the Treasury, Canberra.

Australian Government (2017). Budget Paper. Technical Report 1. URL http://www.budget.gov.au (Accessed 2017-05-21).

Australian Government Department of Veterans' Affairs (2016). Rebalanced assets test to apply from 2017. URL http://www.dva.gov.au/rebalanced-assets-test-apply-2017/ (Accessed 19/10/2016).

Australian Taxation Office (2016). Minimum annual payments for super income streams. URL https://www.ato.gov.au/rates/key-superannuation-rates-and-thresholds/ (Accessed 27/10/2016).

Australian Taxation Office (2017). Super changes. URL https://www.ato.gov.au/individuals/super/super-changes/ (Accessed 03/05/2017).

Avanzi, B. (2010). What is it that makes the Swiss annuitise ? A description of the Swiss retirement system. *Australian Actuarial Journal 16*(2), 135–162.

Avanzi, B. and S. Purcal (2014). Annuitisation and cross-subsidies in a two-tiered retirement saving system. *Annals of Actuarial Science 8*(02), 234–252.

Barillas, F. and J. Fernández-Villaverde (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control 31*(8), 2698–2712.

Baser, O. (2007). Modeling Transformed Health Care Cost with Unknown Heteroskedasticity. *Applied Economics Research Bulletin 1*, 1–6.

Bateman, H. and S. Thorp (2008). Choices and Constraints over Retirement Income Streams: Comparing Rules and Regulations. *Economic Record 84*, 17–31.

Bateman, H., S. Thorp, and G. Kingston (2007). Financial engineering for Australian annuitants. In H. Bateman (Ed.), *Retirement Provision in Scary Markets*, pp. 123–144. Northampton, Massachusetts, USA: Edward Elgar Publishing.

Bäuerle, N. and U. Rieder (2011). *Markov Decision Processes with Applications to Finance* (1 ed.). Universitext. Berlin, Germany: Springer-Verlag Berlin Heidelberg.

Bellman, R. (2003). *Dynamic Programming* (Reprint ed.). New York: Dover Publications.

Bernicke, T. (2005). Reality Retirement Planning: A New Paradigm for an Old Science. *Journal of Financial Planning* (June), 1–8.

Blake, D., D. Wright, and Y. Zhang (2014). Age-dependent investing: Optimal funding and investment strategies in defined contribution pension plans when members are rational life cycle financial planners. *Journal of Economic Dynamics and Control 38*(1), 105–124.

Brandt, M. W., A. Goyal, P. Santa-Clara, and J. R. Stroud (2005). A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability. *Review of Financial Studies 18*(3), 831–873.

Broadie, M. and P. Glasserman (2004). A stochastic mesh method for pricing high-dimensional American options. *Journal of Computational Finance 7*(4), 35–72.

Bütler, M., K. Peijnenburg, and S. Staubli (2016). How much do means-tested benefits reduce the demand for annuities? *Journal of Pension Economics and Finance*, 1–31.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *91*, 312–320.

Chen, H., S. H. Cox, and S. S. Wang (2010). Is the Home Equity Conversion Mortgage in the United States sustainable? Evidence from pricing mortgage insurance premiums and non-recourse provisions using the conditional Esscher transform. *Insurance: Mathematics and Economics 46*(2), 371–384.

Chen, W., N. Langrené, and T. Tarnopolskaya (2015). Switching Surfaces for Optimal Natural Resource Extraction under Uncertainty. In T. Weber, M. McPhee, and R. Anderssen (Eds.), *MODSIM2015, 21st International Congress on Modelling and Simulation*, pp. 1063–1069. Modelling and Simulation Society of Australia and New Zealand. URL `www.mssanz.org.au/modsim2015/E6/chen.pdf` (Accessed 2016-12-02).

Chiang, S. L. and M. S. Tsai (2016). Analyzing an elder's desire for a reverse mortgage using an economic model that considers house bequest motivation, random death time and stochastic house price. *International Review of Economics and Finance 42*, 202–219.

Cho, S. and R. Sane (2013). Means-Tested Age Pensions and Homeownership: Is there a link? *Macroeconomic Dynamics 17*(6), 1281–1310.

Clare, R. (2014). Spending patterns of older retirees: New ASFA Retirement Standard. Technical report, ASFA Research and Resource Centre.

Cocco, J. F. and F. J. Gomes (2012). Longevity risk, retirement savings, and financial innovation. *Journal of Financial Economics 103*(3), 507–529.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*(2), 491–533.

Davidoff, T., J. R. Brown, and P. A. Diamond (2005). Annuities and Individual Welfare. *American Economic Review 95*(5), 1573–1590.

De Nardi, M., E. French, and J. B. Jones (2010). Why Do the Elderly Save? The Role of Medical Expenses. *Journal of Political Economy 118*(1), 39–75.

De Nardi, M. C. (2004). Wealth inequality and intergenerational links. *Review of Economic Studies 71*(3), 743–768.

Denault, M., E. Delage, and J-G. Simonato (2017). Dynamic portfolio choice: a simulation-and-regression approach. *Optimization and Engineering 9*, 1–38.

Denault, M., J-G. Simonato, and L. Stentoft (2013). A simulation-and-regression approach for stochastic dynamic programs with endogenous state variables. *Computers and Operations Research 40*(11), 2760–2769.

Department of Social Services (2014). DSS Demographics June 2014. URL `https://www.data.gov.au/dataset/dss-payment-demographic-data/` (Accessed 2016-11-14).

Department of Social Services (2016). Guides to Social Policy Law. URL `http://guides.dss.gov.au/guide-social-security-law` (Accessed 2016-01-04).

Ding, J. (2014). *Essays on post-retirement financial planning and pension policy modelling in Australia*. PhD dissertation, Macquarie University, Sydney, Australia.

Duan, N. (1983). Smearing estimate: A Nonparametric retransformation method. *Journal of the American Statistical Association 78*(383), 605–610.

Dushi, I. and A. Webb (2004). Household annuitization decisions: simulations and empirical analyses. *Journal of Pension Economics and Finance 3*(2), 109–143.

Fabozzi, F. J., T. Paletta, and R. Tunaru (2017). An improved least squares Monte Carlo valuation method based on heteroscedasticity. *European Journal of Operational Research 263*(2), 698–706.

Fernández-Villaverde, J. and D. Krueger (2007). Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data. *Review of Economics and Statistics 89*(3), 552–565.

Fisher, I. (1930). *The Theory of Interest: As Determined by Impatience to Spend Income and Opportunity to Invest it* (1 ed.). New York, USA: The Macmillan Company.

Friedman, B. M. and M. J. Warshawsky (1990). The Cost of Annuities: Implications for Saving Behavior and Bequests. *The Quarterly Journal of Economics 1*(105), 135–154.

Garlappi, L. and G. Skoulakis (2010). Solving Consumption and Portfolio Choice Problems: The State Variable Decomposition Method. *Review of Financial Studies 23*(9), 3346–3400.

Glasserman, P. and B. Yu (2004). Simulation for American Options: Regression Now or Regression Later? In *Monte Carlo and Quasi-Monte Carlo Methods 2002*, pp. 213–226. Berlin, Heidelberg: Springer Berlin Heidelberg.

Greene, W. (2008). *Econometric Analysis* (6th ed.). New Jersey: Pearson Prentice Hall.

Harvey, A. (1976). Estimating regression models with multiplicative heteroscedasticity. *Econometrica 44*, 461–465.

Henry, K. (2009). Australia' s future tax system - Report to the Treasurer (Overview). Technical Report December, Commonwealth of Australia.

Higgins, T. and S. Roberts (2011). Variability in expenditure preferences among elderly Australians. *Australian National University, presented at the 19th Annual Colloquium of Superannuation Researchers*, 1–28.

Hinz, J. (2014). Optimal Stochastic Switching under Convexity Assumptions. *SIAM Journal on Control and Optimization 52*(1), 164–188.

Hinz, J. and J. Yee (2017). Stochastic switching for partially observable dynamics and optimal asset allocation. *International Journal of Control 90*(3), 553–565.

Hoff, P. D. and X. Niu (2012). A Covariance Regression Model. *Statistica Sinica 22*(2), 1–31.

Huang, H., M. A. Milevsky, and T. S. Salisbury (2012). Optimal retirement consumption with a stochastic force of mortality. *Insurance: Mathematics and Economics 51*(2), 282–291.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995). Precautionary Saving and Social Insurance. *Journal of Political Economy 103*(2), 360–399.

Hull, J. C. (2012). *Options, Futures and Other Derivatives* (8th ed.). Pearson Higher Ed USA.

Hulley, H., R. McKibbin, A. Pedersen, and S. Thorp (2013). Means-Tested Public Pensions, Portfolio Choice and Decumulation in Retirement. *Economic Record 89*(284), 31–51.

Hurd, M. and J. Smith (2003). Expected Bequests and Their Distribution. Technical report, RAND Corporation, Santa Monica, CA.

Hurst, E. and J. Ziliak (2004). Do Welfare Asset Limits Affect Household Saving? Evidence from Welfare Reform. NBER Working Paper No 10487.

Inkmann, J., P. Lopes, and A. Michaelides (2011). How Deep Is the Annuity Market Participation Puzzle? *Review of Financial Studies 24*(1), 279–319.

Iskhakov, F. (2015). Multidimensional endogenous gridpoint method: Solving triangular dynamic stochastic optimization problems without root-finding operations. *Economics Letters 135*, 72–76.

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

Kahaner, D., C. B. Moler, S. Nash, and G. E. Forsythe (1989). *Numerical methods and software* (1 ed.). Englewood Cliffs, New Jersey, USA: Prentice Hall.

Kahneman, D. and A. Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica 47*(2), 263–292.

Karatzas, I., J. P. Lehoczky, S. P. Sethi, and S. E. Shreve (1986). Explicit Solution of a General Consumption/Investment Problem. *Mathematics of Operations Research 11*(2), 261–294.

Kharroubi, I., N. Langrené, and H. Pham (2014). A numerical algorithm for fully nonlinear HJB equations: an approach by control randomization. *Monte Carlo Methods and Applications 20*(2), 145–165.

Kharroubi, I., N. Langrené, and H. Pham (2015). Discrete time approximation of fully nonlinear HJB equations via BSDEs with nonpositive jumps. *The Annals of Applied Probability 25*(4), 2301–2338.

Kingston, G. and L. Fisher (2014). Down the Retirement Risk Zone with Gun and Camera. *Economic Papers 33*(2), 153–162.

Kingston, G. and S. Thorp (2005). Annuitization and asset allocation with HARA utility. *Journal of Pension Economics and Finance 4*(03), 225.

Kou, S., X. Peng, and X. Xu (2016). EM Algorithm and Stochastic Control in Economics. *arXiv:1611.01767*, 1–46.

KPMG (2010). KPMG Econtech CGE Analysis of the Current Australian Tax System. Technical Report March. Available at `https://www.cpaaustralia.com.au`.

Kudrna, G. (2014). Means Testing of Australia's Age Pension: A Numerical Analysis with an OLG Model. Technical report, Centre of Excellence in Population Ageing Research (CEPAR).

Kudrna, G. and A. Woodland (2011a). An inter-temporal general equilibrium analysis of the Australian age pension means test. *Journal of Macroeconomics 33*(1), 61–79.

Kudrna, G. and A. Woodland (2011b). Implications of the 2009 Age Pension Reform in Australia: A Dynamic General Equilibrium Analysis. *Economic Record 87*(277), 183–201.

Kudrna, G. and A. D. Woodland (2013). Macroeconomic and Welfare Effects of the 2010 Changes to Mandatory Superannuation. *Economic Record 89*(287), 445–468.

Langrene, N., T. Tarnopolskaya, W. Chen, Z. Zhu, and M. Cooksey (2015). New Regression Monte Carlo Methods for High-dimensional Real Options Problems in Minerals industry. *21st International Congress on Modelling and Simulation*, 1077–1083.

Lim, B. H., Y. H. Shin, and U. J. Choi (2008). Optimal investment, consumption and retirement choice problem with disutility and subsistence consumption constraints. *Journal of Mathematical Analysis and Applications 345*(1), 109–122.

Lim-Applegate, H., P. McLean, P. Lindenmayer, and B. Wallace (2006). New age pensioners trends in wealth. *Australian Social Policy*, 1–26.

Lockwood, L. M. (2014). Incidental Bequests: Bequest Motives and the Choice to Self-Insure Late-Life Risks. NBER Working Paper No 20745.

Longstaff, F. A. and E. S. Schwartz (2001). Valuing American Options by Simulation: A Simple Least-Squares Approach. *Review of Financial Studies 14*(1), 113–147.

Lucas, R. E. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy 1*, 19–46.

Luo, X. and P. V. Shevchenko (2014). Fast and simple method for pricing exotic options using GaussHermite quadrature on a cubic spline interpolation. *Journal of Financial Engineering 1*(4), 1450033:1–1450033:31.

Luo, X. and P. V. Shevchenko (2015). Valuation of variable annuities with guaranteed minimum withdrawal and death benefits via stochastic control optimization. *Insurance: Mathematics and Economics 62*, 5–15.

Marín-Solano, J. and J. Navas (2010). Consumption and portfolio rules for time-inconsistent investors. *European Journal of Operational Research 201*(3), 860–872.

Merton, R. (1969). Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case. *Review of Economics and Statistics 51*(3), 247–257.

Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory 3*(4), 373–413.

Milevsky, M. and H. Huang (2011). Risk Management Spending Retirement on Planet Vulcan: The Impact of Longevity Risk Aversion on Optimal Withdrawal Rates. *Risk Management* (23), 24–38.

Milevsky, M. A. and V. R. Young (2007). Annuitization and asset allocation. *Journal of Economic Dynamics and Control 31*(9), 3138–3177.

Modigliani, F. and R. Brumberg (1954). Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data. In K. Kurihara (Ed.), *Post Keynesian Economics*, pp. 388–436.

Muller, H.-G. and U. Stadtmuller (1987). Estimation of Heteroscedasticity in Regression Analysis. *The Annals of Statistics 15*(2), 610–625.

Nadarajah, S., F. Margot, and N. Secomandi (2017). Comparison of least squares Monte Carlo methods with applications to energy real options. *European Journal of Operational Research 256*(1), 196–204.

Nadarajah, S. and N. Secomandi (2017). Relationship between least squares Monte Carlo and approximate linear programming. *Operations Research Letters 45*(5), 409–414.

Nakajima, M. (2017). Reverse Mortgage Loans: A Quantitative Analysis. *Journal of Finance 72*(2), 911–950.

Neumann, J. V. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

Neumark, D. and E. Powers (1998). The effect of means-tested income support for the elderly on pre-retirement saving: evidence from the SSI program in the U.S. *Journal of Public Economics 68*(2), 181–206.

Novikov, A. A. and A. N. Shiryaev (2005). On an Effective Solution of the Optimal Stopping Problem for Random Walks. *Theory of Probability & Its Applications 49*(2), 344–354.

Olsberg, D. and M. Winters (2005). Ageing in place: intergenerational and intrafamilial housing transfers and shifts in later life. Technical Report 88, Australian Housing and Urban Research Institute, Melbourne.

Plan For Life (2016). Report on the Australian Reitrement Income Market. Technical report, Plan For Life. Available at `http://www.pflresearch.com.au/`.

Poterba, J., J. Rauh, S. Venti, and D. Wise (2007). Defined contribution plans, defined benefit plans, and the accumulation of retirement wealth. *Journal of Public Economics 91*(10), 2062–2086.

Rice Warner (2015). Quo Vadis ? Superannuation needs effective policy... not politics. *Submission to Tax White Paper Task Force*.

Richard, S. F. (1975). Optimal consumption, portfolio and life insurance rules for an uncertain lived individual in a continuous time model. *Journal of Financial Economics 2*(2), 187–203.

Rothman, G. (2012). Modelling the sustainability of Australia's retirement income system. *Department of Treasury, presented to the 20th Colloqium of Superannuation Researchers*.

Samuelson, P. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics 51*(3), 239–246.

Shao, A. W., K. Hanewald, and M. Sherris (2015). Reverse mortgage pricing and risk analysis allowing for idiosyncratic house price risk and longevity risk. *Insurance: Mathematics and Economics 63*, 76–90.

Shevchenko, P. V. (2016). Analysis of withdrawals from self-managed super funds using Australian Taxation Office data. *CSIRO Technical Report EP164438*. CSIRO Data61.

Shin, Y. H. (2012). Voluntary retirement and portfolio selection: Dynamic programming approaches. *Applied Mathematics Letters 25*(7), 1087–1093.

Shin, Y. H., B. H. Lim, and U. J. Choi (2007). Optimal consumption and portfolio selection problem with downside consumption constraints. *Applied Mathematics and Computation 188*(2), 1801–1811.

Smyth, G. K. (1989). Generalized linear models with varying dispersion. *Journal of the Royal Statistical Society 51*(1), 47–60.

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

Sun, W., R. K. Triest, and A. Webb (2008). Optimal Retirement Asset Decumulation Strategies: The Impact of Housing Wealth. *Asia-Pacific Journal of Risk and Insurance 3*(1), 123–149.

The Commonwealth of Australia (2015). Budget 2015 Overview. Technical report. URL `http://www.budget.gov.au` (Accessed 2014-10-01).

Tran, C. and A. Woodland (2014). Trade-offs in means tested pension design. *Journal of Economic Dynamics and Control 47*, 72–93.

Tsitsiklis, J. N. and B. Van Roy (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks 12*(4), 694–703.

World Bank (2008). The World Bank Pension Conceptual Framework. (202), 8.

Yaari, M. (1964). On the consumer's lifetime allocation process. *International economic review 5*(3).

Yaari, M. (1965). Uncertain Lifetime, Life Insurance, and the Theory of the Consumer. *The Review of Economic Studies 32*(2), 1–137.

Yao, H., Y. Lai, Q. Ma, and M. Jian (2014). Asset allocation for a DC pension fund with stochastic income and mortality risk: A multi-period mean-variance framework. *Insurance: Mathematics and Economics 54*(1), 84–92.

Yogo, M. (2016). Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets. *Journal of Monetary Economics 80*, 17–34.

Zhang, R., N. Langren, Y. Tian, Z. Zhu, F. Klebaner, and K. Hamza (2016). Dynamic Portfolio Optimisation with Intermediate Costs: A Least-Squares Monte Carlo Simulation Approach. *Preprint SSRN: 2696968*. Available at `https://papers.ssrn.com`.

Zhang, R., N. Langrené, Y. Tian, Z. Zhu, F. Klebaner, and K. Hamza (2016). Efficient Simulation Method for Dynamic Portfolio Selection with Transaction Cost, Liquidity Cost and Market Impact. *Preprint arXiv: 1610.07694*. Available at `https://arxiv.org`.

Zhou, X., H. Lin, and E. Johnson (2008). Nonparametric Heteroscedastic Transformation Regression Models for Skewed Data with an Application to Health Care Costs. *Journal of the Royal Statistical Society. Series B (Statistical Methodology) 70*(5), 1029–1047.