DEPARTMENT
OF ECONOMICS

Essays in Macroeconomics: Investment Choices, Labor Market Outcomes, and Family Structure

Annika Bacher

Thesis submitted for assessment with a view to
obtaining the degree of Doctor of Economics
of the European University Institute

Florence, 30 May 2022

European University Institute
Department of Economics

Essays in Macroeconomics: Investment Choices, Labor Market Outcomes, and Family Structure

Annika Bacher

Thesis submitted for assessment with a view to
obtaining the degree of Doctor of Economics
of the European University Institute

Examining Board

Prof. Russell Cooper, EUI, Supervisor
Prof. Thomas Crossley, EUI, Co-Supervisor
Prof. Dean Corbae, University of Wisconsin-Madison
Prof. Michael Haliassos, Goethe University Frankfurt

Researcher declaration to accompany the submission of written work Department Economics - Doctoral Programme

I Annika Bacher certify that I am the author of the work "Essays in Macroeconomics: Investment Choices, Labor Market Outcomes, and Family Structure" I have presented for examination for the Ph.D. at the European University Institute. I also certify that this is solely my own original work, other than where I have clearly indicated, in this declaration and in the thesis, that it is the work of others.

I warrant that I have obtained all the permissions required for using any material from other copyrighted publications.

I certify that this work complies with the Code of Ethics in Academic Research issued by the European University Institute (IUE 332/2/10 (CA 297).

The copyright of this work rests with its author. Quotation from it is permitted, provided that full acknowledgement is made. This work may not be reproduced without my prior written consent. This authorisation does not, to the best of my knowledge, infringe the rights of any third party.

I declare that this work consists of 58,715 words.

Statement of inclusion of previous work:

I confirm that chapter 3 was jointly co-authored with Philipp Grübener and Lukas Nord and I contributed 33% of the work.

Signature and date:

A. Bacher

12.05.2022

Abstract

This thesis is composed of three independent essays in heterogeneous agent macroeconomics. They all explore how family structure affects investment choices and labor market outcomes of individuals.

The first chapter, **Housing and Savings Behavior Across Family Types**, studies how family structure in the form of marital status and changes thereof affect housing demand. I develop a life-cycle model of housing and financial portfolio choice with dynamic and heterogeneous family types that I calibrate to household survey data from the United States. My findings indicate that divorce risk encourages precautionary savings of couples and reduces their demand risky assets and for indivisible housing. Prospective marriage, lower income levels and larger exposure to income fluctuations prevent singles from becoming homeowners. As a result, abstracting from distinct family types overstates the effectiveness of housing policies such as lowering property taxes and reducing housing transaction costs by up to over 100%. This misspecification is largest among young households, who are most likely to be single and whose marital transition risk is highest. In contrast, regulations that facilitate stock market participation help to foster wealth accumulation, because they encourage investment in high return assets that are cheaper to liquidate in the event of a marital or labor income shock.

In the second chapter, **The Gender Investment Gap over the Life-Cycle**, I document with data from the Survey of Consumer Finances that single women hold on average less risky portfolios than single men. To understand the sources of this "Gender Investment Gap", I develop a life-cycle model of portfolio choice that allows for gender differences along observable characteristics and stochastic processes. The model is able to replicate the empirical patterns without introducing gender heterogeneity in preferences. Counterfactual simulations reveal that lower income levels and larger household sizes (mainly through the presence of children) of single women make it optimal for them to invest less risky. Hence, the gender wage gap gets amplified because it results in investment behavior that pays on average lower returns. Importantly, not only current-period income levels and number of household members help to explain this finding but also expectations over their future realizations.

The third chapter, **Joint Search over the Life-Cycle**, co-authored with Philipp Grübener and Lukas Nord, focuses on labor market outcomes and couples. Specifically, we study how intra-household insurance against individual job loss through increased spousal labor market participation, also called the added worker effect, varies over the life cycle. First, we show in U.S. data that the added worker effect is much stronger for young than for old households. A stochastic life cycle model of two-member households with job search in a frictional labor market is capable of replicating this finding. The model suggests that a lower added worker effect for the old is driven primarily by better insurance through asset holdings. Human capital differences between employed young and old contribute to the difference but are quantitatively less important, while differences in job arrival rates play a limited role.

Acknowledgements

This thesis marks the end of an incredibly challenging but also rewarding journey which would not have been possible without the support of many people.

First and foremost, I would like to thank Russell Cooper for his constant guidance and support, his ability to ask the right questions and to encourage me to believe in myself and in my research. My gratitude also goes to Tom Crossley for his enduring support and encouragement. He somehow always managed to ask different questions than Russell. Together, they formed an invaluable team of support for me. I would also like to thank Michael Haliassos and Dean Corbae for generously agreeing to be part of my thesis committee.

Moreover, I am hugely indebted to Mariacristina de Nardi. I met her during my second PhD year, and she has been a constant role of support since. From inviting me to visit her in Minneapolis, to letting me participate in her workshop, and for all her help and encouragement during the job market. I cannot express in words how much I learned from these experiences. Many thanks also Jesús Bueren to whom I could always turn to for advice, feedback, and support.

I owe special thanks to my friends Philipp Grübener and Lukas Nord (who also happen to be my co-authors), from whom I learn every day and who not only contributed immensely to the quality of this thesis but also to my joy of writing it. I am grateful to have been part of a particular great cohort at the EUI who always supported each other and with whom I had so much fun during the past years. Especially, I thank the "German boys" Jojo, Lukas and Wolfram who made me feel home in Florence right from the beginning.

Lastly, I would like to thank Johannes, Maren, and my parents who make my life as great as it is. They were there for me along every step of the way and unconditionally support me in everything I do. You mean the world to me and I will be forever grateful to have you in my life.

Contents

Chapter 1

Housing and Savings Behavior Across Family Types

Abstract Does marital status affect households' investment choices? Is accounting for distinct family types necessary for the correct evaluation of policies that aim at stimulating housing demand? To answer these questions, I develop a life-cycle model of housing and financial portfolio choice with dynamic and heterogeneous family types. I find that divorce risk encourages precautionary savings of couples in the form of liquid assets and reduces their demand for illiquid housing. Expected marriage, low income levels and larger exposure to income fluctuations prevent singles from becoming homeowners. Abstracting from distinct family types amplifies the attractiveness of housing and, as a result, overstates the effectiveness of housing policies such as lowering property taxes and reducing transaction costs by a factor greater than two. This mis-specification is largest for young households who are most likely to be single and whose marital transition risk is highest. In contrast, regulations that facilitate stock market participation help to foster wealth accumulation, because they encourage investment in high return assets that are cheaper to liquidate in the event of a (marital or labor income) shock.

1.1 Introduction

In the United States, housing represents the largest asset in most households' portfolios and constitutes the primary way through which they accumulate wealth (Goetzmann et al., 2021). In addition, being a homeowner is often regarded as part of the American Dream, and has been shown to improve children's outcomes and to strengthen involvement with the local community.[1] As a result, increasing homeownership attracts considerable attention among policy makers and has resulted in numerous policy proposals targeted at stimulating housing demand.[2] Today, the United States alone invests around $200 billion annually to finance policies that promote homeownership (Sodini et al., 2021).

However, to evaluate the transmission of any policy that operates through housing demand, it is necessary to understand the determinants of households' investment choices over their life-cycle. The literature has so far identified a variety of household characteristics that shape the demand for housing. Examples of these include age, income dynamics, or wealth holdings (e.g. Attanasio et al., 2012; Paz-Pardo, 2020). In this paper, I argue that marital status is another important, yet understudied, driver of households' housing and investment decisions because it affects labor income profiles and because their illiquidity makes houses expensive to liquidate in the event of marriage or divorce.

In particular, I address two questions: First, what are the key channels through which marital status affects investment dynamics of couples vs. singles? Second, does accounting for distinct family types change the effectiveness of housing policies over the life-cycle?

I first document novel empirical patterns on heterogeneity in financial portfolio composition and in housing choices across couples, single men and single women in the United States by combining data from the Survey of Consumer Finances (SCF) and from the Panel Study of Income Dynamics (PSID). I show that on average almost 80% of all couples own their house, whereas less than half of all single households do. In contrast, the average house value of couple owners is per capita $55,000 lower than that of single men and $29,000 lower than that of single women. Moreover, at retirement age, the average couple household has accumulated per capita around $50,000 more in financial savings than the average single man, and around $150,000 more than the average single woman.

Next, I develop a life-cycle framework of housing, financial portfolio choice and family structure that is able to replicate these empirical patterns. In the model, households derive utility from nondurable consumption and from housing services. They decide on consumption, savings in safe and in risky financial assets as well as their housing stock, forming expectations about future labor income, asset returns and marital status. Housing is discrete, giving rise to a minimum house size available for

[1] See Forbes (2019) or Goodman and Mayer (2018) on housing and the American Dream, Haurin et al. (2002) on children's outcomes, and DiPasquale and Glaeser (1999) on homeownership and local community involvement.

[2] For example, President Biden declared June 2021 as "National Homeownership Month" and explicitly called "...to recognize the enduring value of homeownership and recommit ourselves to helping more Americans realize that dream". Policy examples include housing subsidies to first-time buyers, tax credits, the home mortgage interest deduction (HMID) but also reforms that aim at reducing property or transfer taxes.

purchase. In addition, housing adjustments are subject to transaction costs and homeowners have to pay annual maintenance costs.

Family types are heterogeneous in terms of their labor income profiles which I estimate separately for single women, single men and couples from the PSID. I find that couple households have on average higher labor income levels than singles. At the same time, they are exposed to smaller labor income fluctuations which in turn affects their willingness to bear risk along other dimensions, for example in financial markets (Heaton and Lucas, 2000).

Additionally, couples face the possibility of getting divorced whereas singles may meet a partner whom they marry. Both events impose substantial financial uncertainty that works in opposite directions. Divorce constitutes a financial risk because it requires households to split their assets and because it results in a state with lower labor income levels and higher labor income risk. Marriage, in contrast, crowds out private savings as it reflects a financial outcome with disproportionally high returns (through asset holdings of the partner) and the ability to pool income within the household. In order to realistically replicate this financial uncertainty, I require the model to match empirical shifts in homeownership rates and in financial wealth throughout the years preceding and following a marital shock.

Moreover, households enjoy economies of scale which differ between housing and non-durable consumption, to capture that housing services might be more easily divided among family members than non-durable consumption items (Yang, 2009). Hence, heterogeneity in the number of household members by age and by family type affects both the optimal allocation of resources across time and the optimal intratemporal allocation across goods.

Key Channels. I calibrate the model to match key moments on ownership, savings behavior, stock market participation and house prices in the US. By means of counterfactual simulations, I show that marital transition risk and heterogeneity in labor income profiles are the most important determinants through which marital status affects individual housing demand.

The income and asset losses associated with divorce induce couples to increase precautionary savings in safe assets and lower their demand for risky assets and for illiquid housing. On aggregate, the increased savings motive dominates the portfolio shift towards low-return assets, resulting in larger wealth holdings than in a world without divorce. In addition, the increased savings motive shifts the distribution of couples in the economy towards asset-richer households who are more likely to be homeowners. Therefore, allowing for divorce increases the aggregate share of owning couples, despite lower housing demand.

In contrast, marriage represents an outcome with disproportionately high returns through asset holdings of the partner. Furthermore, it increases households' prospective savings ability because of higher income levels and resource sharing, crowding out private savings of singles relative to a world without marital transitions. At the same time, marriage reflects an event that may render a previously purchased home suboptimal, reducing singles' housing demand. Consequently, both aggregate financial savings and homeownership rates of singles decrease when allowing for marriage in the model.

Low labor income levels further prevent singles from accumulating financial savings and keep them out of homeownership. The corresponding drop in household income following a divorce strengthens couples' savings motive, additionally contributing to the empirical observation that couples are more likely to be homeowners and that they hold more financial wealth than singles.

Similary, larger labor income fluctuations of single households increase their own risk exposure but also that of couples in the event of divorce. Thus, more labor market risk of singles reduces the demand for risky assets (to reduce overall risk exposure) and for housing (which cannot be easily liquidated in the event of an adverse labor income shock) of all households. For singles, this mechanism operates most strongly along the extensive margin (that is, households shift from small owner-occupied housing towards renting). As a result, the average housing wealth of owning singles is tilted towards larger homes, contributing to the empirical pattern that, conditional on owning, singles invest more wealth into housing than couples.

Implications for Policy Evaluation. Using the calibrated model, I show that accounting for family composition is quantitatively important for the evaluation of policies that aim at stimulating housing demand. I simulate two types of reforms: lowering housing transaction costs and reducing property taxes. Thus, the first policy facilitates housing adjustments in response to shocks whereas the latter lowers the flow costs of housing. I then perform the same exercises in a standard framework with one generic household type and compare the effectiveness of both reforms in terms of increasing homeownership rates across set-ups.

My main results are as follows. Introducing marriage and divorce lowers the attractiveness of indivisible housing and aggregate homeownership rates increase less in response to both policies. Additionally, the presence of single households who have low income levels and who are exposed more to labor income fluctuations (and are thus less likely to invest into housing) further weaken policy transmission. Quantitatively, abstracting from distinct family types overstates the effectiveness of lowering property taxes by 133% and that of decreasing transaction costs by 53%. Hence, the framework with one generic household type not only overestimates the effectiveness of both reforms and but also biases their relative magnitude. Lowering property taxes shields households from large expenditure commitments each period. However, once I allow for marital transitions, households value relatively more to be able to adjust their housing size (in the event of a marital shock) at little cost.

In addition, because marriage and divorce probabilities are decreasing in age, the magnitude of this mis-specification across frameworks is largest for young households. Since young households are the age group that most housing policies in the US are primarily targeted at, this result further emphasizes the importance of accounting for distinct family types when designing or evaluating reforms that aim at stimulating housing demand.[3]

Lastly, I evaluate both policies in terms of increasing households' net worth and as an alternative consider a regulation that facilitates stock market participation. I find that the latter is most effective in fostering overall wealth accumulation, especially among singles, since it encourages investment in assets that pay high returns in expectations but allow for relatively small investment amounts and can be more easily pooled in the event of marriage.

Related Literature. This paper contributes to several strands of the literature. Broadly, it relates to a large literature on housing in macroeconomics and on financial portfolio allocation of households. Piazzesi and Schneider (2016) provide a detailed review of the former and Gomes et al. (2020) as well as Campbell (2006) of the latter. For a literature review on life-cycle dynamics of household's portfolio composition, see Poterba and Samwick (2001) and, more recently, Gomes (2020).

More specifically, I complement a literature that studies the interaction of housing dynamics and a financial portfolio choice within life-cycle frameworks (Cocco (2005), Yao and Zhang (2005), Flavin and Yamashita (2011), Chetty et al. (2017), Vestman (2019), Paz-Pardo (2020), Brandsaas (2021)). Expanding on their work, I am the first to introduce distinct family types and am thus able to quantify the importance of marital status on housing decisions as well as their interaction with a financial portfolio choice.

Furthermore, my paper extends previous work that analyzes how marital dynamics affect household investment decisions. For example, Fisher and Gervais (2011), Fischer and Khorunzhina (2018), Chang (2019), Khorunzhina and Miller (2019), and Bartscher (2020) study the effect of marriage and divorce on home-buying decisions and mortgage applications. Love (2010) and Hubener et al. (2015) develop a

[3] Typically, most housing policies are targeted at first-time buyers with the explicit goal of stimulating housing demand among young ("millennial") households. See for example Choi et al. (2018).

joint framework of household structure and financial portfolio choice that abstracts from housing to study how men and women re-balance their financial portfolio following family shocks such as divorce. More generally, many papers focus on the interaction of marital transition dynamics or marital status and household savings behavior (e.g. Cubeddu and Ríos-Rull (2003), Yamaguchi et al. (2014), Voena (2015), Fehr et al. (2016), Borella et al. (2018), and De Nardi et al. (2021b)). Some empirical work such as Stevenson (2007), Mundra and Uwaifo Oyelere (2016), and Goodman et al. (2019) investigates the determinants of housing choices conditional on marital status. Relative to these papers, my focus is on how marital status interacts with housing demand over the life-cycle and on deriving implications for policies that aim at stimulating homeownership.

My paper is closest to Peter et al. (2020) who propose a joint framework of housing choices and marital status to study homeownership rates between singles and couples across Europe. Their findings indicate that higher homeownership rates among couples can be attributed to weak rental markets or strong credit markets, depending on the specific country under consideration. In contrast to their work, my paper is limited to one country (the US), additionally includes a financial portfolio choice between safe and risky assets and emphasizes life-cycle dynamics of portfolio composition depending on the family type.

Finally, my paper is related to a macroeconomic literature on life-cycle dynamics of portfolio composition with durable goods. Attanasio et al. (2012) study the channels through which housing demand evolves over the life-cycle. Fernández-Villaverde and Krueger (2011) emphasizes the importance of housing as collateral because it relaxes borrowing constraints, explaining why households accumulate housing early in life and only later start saving in financial assets. Albeit being present in my framework as well, this channel is weakened through the introduction of single households who are reluctant to invest in housing early in life as they expect to get married soon. Similarly, Yang (2009) focuses on life-cycle patterns of consumption and shows that the collateral value of housing is key to replicate the increasing housing stock early in life, while its illiquid nature can account for the slow decumulation of housing among the old. In relation to her, I provide evidence that the illiquidity of housing not only helps to understand its slow decumulation during old age but also the marital gap in homeownership across couples and singles.

Roadmap. The remainder of this paper is structured as follows. Section 1.2 presents empirical evidence on life-cycle patterns of portfolio dynamics of single men, single women and couples. Section 1.3 introduces the structural model. Next, Section 1.4 explains the calibration strategy, Section 1.5 analyzes the channels through which marital status affects household's investment choices and Section 1.6 discusses implications for policy evaluation. Section 1.7 concludes.

1.2 Key Facts

The following section first documents key differences in investment behavior across couples, single men and single women over their life-cycle, relying on data from the Survey of Consumer Finances, waves 1989-2016.[4] Second, to further shed light on how marital risk interacts with households' investment decisions, I conduct an event study of housing and financial portfolio choices around the time of marriage and divorce. Later on, I will validate the performance of the structural model with regard to these empirical patterns.

1.2.1 Life-Cycle Patterns of Investment Choices Across Families

Figure 1.1a shows that the share of homeowners among couples is higher than among both single men and single women at every age in the US. On average, this "marital gap" in homeownership rates is around 30%pts, corresponding to the share of single owners being 46% lower than the share of couple owners. Single women refer to family units with a female head without a spouse. Single men are defined accordingly. Couples include legally married individuals with both spouses present in the household.[5] Figure 1.1b documents the average housing wealth of homeowners across family types. Conditional on owning, couples allocate (per capita) on average $44,000 less wealth into housing than singles.[6] Thus, couples invest more in housing along the extensive margin whereas singles tend to invest more along the intensive margin, once they become owners. Moreover, while I find hardly any gender differences in the share of single owners, the conditional housing wealth of single men is higher than of single women, in particular during older ages.

Figure 1.2 considers financial investment patterns across family types. In contrast to housing wealth, couples accumulate more financial assets (per capita) than both single men and single women (Figure 1.2a).[7] Financial assets are defined as the sum of all risky and safe financial assets that the household

[4] Appendix A.1 describes the data and sample selection criteria in detail.

[5] In Appendix A.2, I show that the reported empirical patterns are robust to including cohabiting individuals either in the "couples" or in the "singles" category.

[6] Figure 1.1b displays the mean conditional house value, irrespectively of any housing debt. Appendix A.1.2 shows that this finding is robust to considering the median as well as to considering housing equity.

[7] Again, this finding is robust to considering median of financial assets, see Appendix A.1.2 for details.

Figure 1.1 Housing Choices Across Family Types (Data)

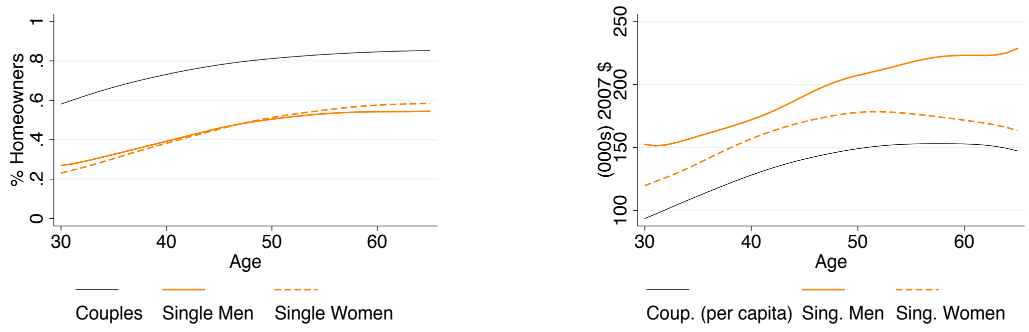

(a) Homeownership Rates **(b)** (Conditional) Mean House Value

Notes: Figure 1.1 plots the life-cycle profiles of homeownership rates and average house value of owners by family type. House value is defined as the value of a household's primary residence, irrespective of any mortgage debt. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

holds. Risky assets refer to direct stock holdings, corporate and foreign bonds, the fraction of mutual funds that is invested in the former as well as the fraction of retirement accounts which is invested in stocks. Safe financial assets include cash holdings, savings and checking accounts, government bonds as well as the fraction of mutual funds and retirement accounts which is invested in safe assets.

At the entry of retirement, the average financial wealth of single women is little over $100,000, that of single men almost $200,000 and couples hold on average per capita $250,000 in financial assets. This pattern prevails when considering only risky financial assets (Figure 1.2b). However, Figure 1.2b combines the extensive margin (i.e. whether or not the household holds any risky assets) and the intensive margin of risky asset holdings. When plotting both margins separately (reported in Appendix A.1.2), I find that, as for housing wealth, couples are more likely to participate in risky asset markets, but that they, conditional on participation, do not hold more risky assets than singles. Single men hold persistently more financial wealth than single women, again with this gender gap widening in age.[8]

Finally, Figure 1.3 plots the portfolio shares of housing, risky financial assets and safe financial assets by family type and by age group. The housing share is defined as housing equity over overall wealth.[9] In line with Figure 1.1, the housing share of couples is higher than that of singles. Additionally, the housing share of couples is relatively flat over their life-cycle whereas that of both single men and single women is increasing in age.

8 Appendix A.1.2 reports the life-cycle profiles of housing and financial wealth accumulation of never married singles vs. divorced individuals.

9 Figure A.4 in Appendix A.1.2 illustrates how these patterns differ when splitting the housing share into mortgages and house value (as opposed to considering housing equity).

Figure 1.2 Financial Choices Across Family Types (Data)

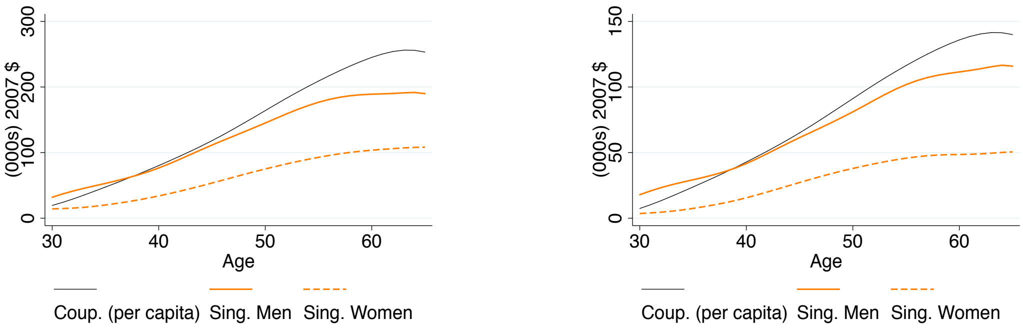

(a) Mean Financial Assets **(b)** Mean Risky Assets

Notes: Figure 1.2 plots average financial assets and average risky assets by family type. Financial assets are defined as the sum of safe and risky financial assets. Risky assets contain direct stock holdings, corporate and foreign bonds, the fraction of mutual funds that include the former as well as the fraction of retirement accounts which is invested in stocks. Safe financial assets refer to cash holdings, savings and checking accounts, government bonds and the fraction of mutual funds and retirement accounts which is invested in safe assets. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

Figure 1.3 Portfolio Shares by Age (Data)

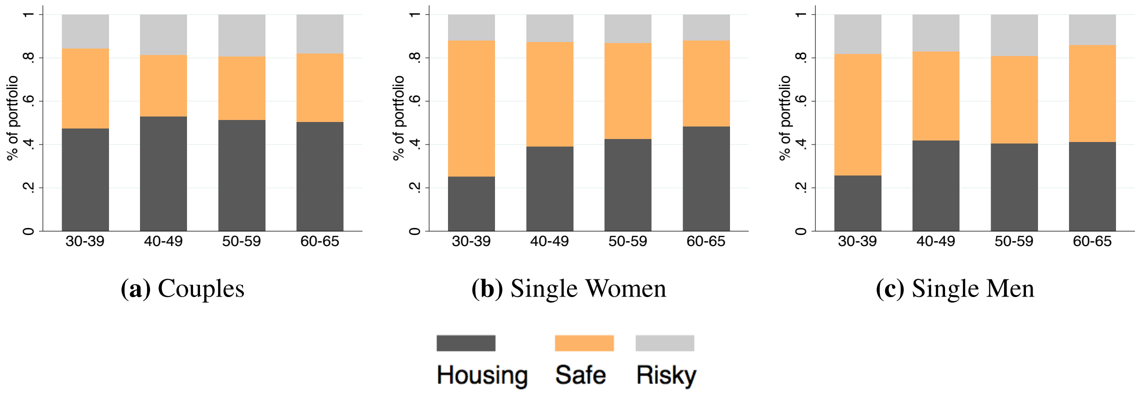

(a) Couples **(b)** Single Women **(c)** Single Men

Notes: Figure 1.3 plots the average share of overall wealth invested in housing, safe and risky assets by family type and age category. The housing share denotes housing equity as a fraction of overall wealth (the sum of the house value, safe and risky assets net of any mortgage debt). Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

1.2.2 Housing and Financial Portfolio Choices around Marital Shifts

To understand how marital risk interacts with investment choices of households, one key dimension is households' housing and financial portfolio shifts around marriage and divorce. These shifts directly affect the financial riskiness of marital transitions and hence, households' overall risk exposure. In this section, I conduct an event study of housing- and asset choices in the years preceding and following marriage and divorce.[10]

Divorce. Figure 1.4 documents average homeownership rates, stock market participation rates and median financial asset holdings around the time of divorce. All values refer to household level estimates.

[10] To do so, I work with data from the PSID because of its panel structure.

The year zero indicates the first year in which the respondent reports to be divorced. Following a divorce (i.e. between year -2 and year 0), homeownership rates drop by around 30% and median financial assets by around 50% because spouses have to split up their assets. In the years after divorce, median financial assets and the share of homeowners gradually increase again. In contrast, the average stock market participation rate (Figure 1.4b) is hardly affected by the separation: it slightly declines in the years prior to divorce and stays mostly flat afterwards.

Figure 1.4 Housing and Financial Portfolio Allocation around Divorce (Data)

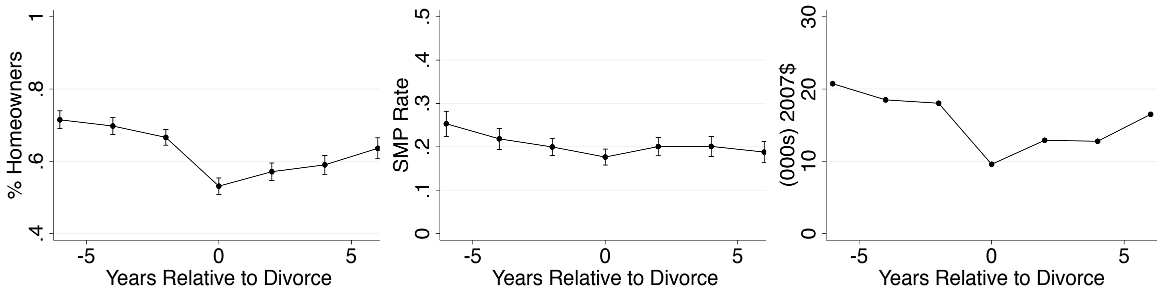

(a) Homeownership Rate **(b)** Stock Market Participation **(c)** Median Financial Assets

Notes: Figure 1.4 plots the the evolution of homeownership rates, stock market participation rates and median financial assets in the years preceding and following a divorce. All values refer to household level estimates. The year zero indicates the first observation in which the individual reports to be divorced. Data is from the Panel Study of Income Dynamics (PSID), waves 1989-2016.

Marriage. Figure 1.5 reports average homeownership rates, stock market participation rates and median financial asset holdings around the time of marriage. Again, all values refer to household level estimates. The year zero indicates the first year in which the respondent reports to be married. After getting married (that is, between year -2 and year 0), the average homeownership rate as well as median financial assets rise continuously. This increase captures both an age effect but also reflects that (newly) married households accumulate more financial wealth than singles and are more likely to become homeowners. The share of stock market participants slightly jumps in the period of marriage and remains flat afterwards (Figure 1.5b).

1.2.3 Robustness Exercises

I conduct several sensitivity checks with respect to the reported marital gaps in investment choices and list the results in Appendix A.2. The documented gaps are robust to replicating the analysis on one cohort of individuals, to including cohabiting households either in the "couples" or "singles" category, and to excluding the housing boom period in the early 2000s as well as the years of the Great Recession, which are both periods with either strong increases or drops in house prices.

Figure 1.5 Housing and Financial Portfolio Allocation around Marriage (Data)

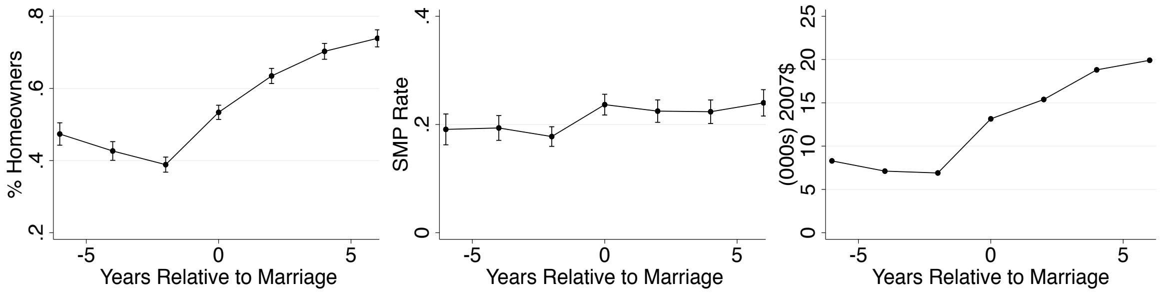

(a) Homeownership Rate **(b)** Stock Market Participation **(c)** Median Financial Assets

Notes: Figure 1.5 plots the evolution of homeownership rates, stock market participation rates and median financial assets in the years preceding and following a marital union. All values refer to household level estimates. The year zero indicates the first observation in which the individual reports to be married. Data is from the Panel Study of Income Dynamics (PSID), waves 1989-2016.

1.3 A Life-Cycle Model of Housing and Portfolio Choice

In this section, I develop a stochastic life-cycle model with three types of households: single men, single women and couples. Time is discrete and the model period is two years. Agents enter the model at age 30, retire deterministically at age 64 and live at most for 84 years, that is $j \in \{30, 32, \ldots, 64, \ldots, 84\}$. Households value non-durable consumption and derive utility from housing. During the working stage, they are subject to idiosyncratic labor income shocks and allocate their portfolio between illiquid housing, liquid safe assets and liquid risky assets. They face marital transition shocks that depend on their current labor income, their age and in the case of marriage, their gender. To purchase a home, households have access to collateralized borrowing (mortgages). During retirement, agents' marital status is fixed, they receive a flat pension payment, and face a positive probability of dying. They can invest in housing, safe and risky assets, and can take out loans in the form of mortgages. At age 84, households have to re-pay all their debt. Upon dying, agents value leaving bequests.

1.3.1 Preferences

Households derive utility from nondurable consumption c and from housing services s. As common in the literature (e.g. Yang, 2009), I express the per-period utility function as:

$$\frac{(g(c,s))^{1-\gamma}}{1-\gamma}$$

where γ denotes the coefficient of relative risk aversion and $g(c,s)$ is specified as:

$$g(c,s) = \left(\omega\left(\frac{c}{\eta_{ij}^{c}}\right)^{\nu} + (1-\omega)\left(\frac{s}{\eta_{ij}^{s}}\right)^{\nu}\right)^{\frac{1}{\nu}}$$

The term ω measures the taste for housing services relative to nondurable consumption goods and ν specifies the substitutability between these two goods. The terms η_{ji}^{c} and η_{ji}^{s} are demographic shifters for changing household sizes over the life-cycle. They vary by age j, household type i (couple, single woman or single man) and are allowed to differ between nondurable consumption goods and housing services to take into account that housing services may be more easily shared than non-durable consumption goods such as food or clothing (Nelson, 1988). The term η is smaller than the overall number of household members and indicates economies of scale. Hence, differences in household size alter the optimal allocation of resources across goods within one period in addition to affecting the optimal allocation of resources over time.

Bequest Motive

In the event of death, individuals derive utility from leaving bequests as in De Nardi (2004):

$$\phi(a', H') = L\frac{(\xi + a' + p_h H')^{1-\gamma}}{1-\gamma}$$

where a' denotes financial assets, $p_h H'$ is the value of the house, ξ captures the luxuriousness of the bequest motive and L governs the bequest intensity. Couples value leaving bequests if they both die within the same period. Whenever only one spouse dies, the surviving spouse keeps the house and continues life as a single with a fraction of the couples' financial assets to account for bequests to non-spousal heirs.

1.3.2 Children

Children enter the model as a deterministic function of age, gender and marital status through changes in the demographic shifters η^c and η^s. In particular, I compute the average number of children by marital status, gender and age from the data and allocate that number of children to all agents in the model who are in the respective age group and have the respective household type. The choice to introduce children in a rather parsimonious way is supported by the data: in Appendix A.1.2, I show that, conditional on family type, investment choices of households (both in terms of homeownership rates as well as in terms of financial wealth accumulation) with and without children are not significantly different from one another. Hence, marital status *per se* seems to be a more important driver of portfolio allocation decisions than the presence of children.

1.3.3 Household Earnings

Working Age. During working age, households supply labor inelastically and face uninsurable income shocks. Labor income can be split into a deterministic and into a stochastic component. Both of these components vary by household type (couples, single men, single women). Income y_{ij} at age j for household type i can be expressed as:

$$y_{ij} = \bar{x}_i \chi_{ij} \tilde{y}_{ij}$$

where $\bar{x}_i$ denotes the constant and χ_{ij} represents an age-specific term. The term $\tilde{y}_{ij}$ captures the stochastic component of labor income.

Guvenen et al. (2021) and De Nardi et al. (2021a) emphasize higher-order moments of the labor income process and show that households' earnings dynamics are characterized by negative skewness and excess kurtosis, both properties that a normally distributed income shock fails to capture.[11] To account for these properties in my set-up, I parameterize $\tilde{y}_{ij}$ as an AR(1) process in logs with innovations drawn from a Gaussian mixture ("GMAR Process"):

$$\tilde{y'} = \rho \tilde{y} + \nu$$

where $\rho \in (0,1]$ captures the persistence of shock ν which is defined as:

$$\nu = \begin{cases} \mathcal{N} \sim (\mu_1, \sigma_1^2) & \text{with probability } p_{\tilde{y}} \\ \mathcal{N} \sim (\mu_2, \sigma_2^2) & \text{with probability } (1 - p_{\tilde{y}}) \end{cases}$$

For small $p_{\tilde{y}}$, negative μ_1, large σ_1^2 and small σ_2^2, this parameterization allows for negative skewness and excess kurtosis. To keep the process stationary, it has to hold that $\mu_2 = \left(\frac{-p_{\tilde{y}}}{1-p_{\tilde{y}}}\right) \mu_1$.

11 Negative skewness implies that more mass of the earnings distribution is concentrated in its left than in its right tail. Excess kurtosis describes heavy tails, meaning that most households experience very small earnings changes, however when hit by a shock, these tend to be quite large.

Retirement. Pension payments are modeled as a fraction of the household's last realized labor income to mimic in a parsimonious way that in the US pension payments are a fraction of an individual's life-time earnings.

1.3.4 Asset Markets

Financial Assets. Households choose between two types of financial assets: one-period safe assets (a_s) and one-period risky assets (a_r). The safe asset pays a time-invariant return r_s. The return of the risky asset is drawn from the distribution $r_r \sim N(\mu_r, \sigma_r^2)$, which is i.i.d and with $\mu_r > r_s$. Following Fagereng et al. (2017), I allow for the possibility of stock market crashes and augment the return of the risky asset by a "disaster" state. That is, with probability $(1 - p_{tail})$ the return is drawn from the above normal distribution and with probability p_{tail} a tail event $r_{tail} < \underline{\mathrm{r}}_r$ materializes. Whenever households choose to invest part of their financial savings into risky assets, they have to pay a per-period lump-sum participation cost S^F to do so.[12] Moreover, homeowners can borrow in one-period mortgages against the value their house, which entails a borrowing premium, i.e. $r_m > r_s$.[13] Additionally, mortgages are subject to an LTV requirement, meaning that the maximum amount of household debt is a fraction ζ_h of the price of its home.

Housing. Households can either be homeowners or renters and have access to houses of discrete sizes:

$$\mathcal{H} = \{R_1, \ldots, R_R, H_1, \ldots, H_H\}$$

where R denotes renting. Both renters and homeowners derive utility from housing services s that are modeled as a correspondence between the size of the house $\mathcal{H}$ and the consumption benefits s derived from it. Owner-occupied housing H can be bought at a fixed price p_H, which deterministically appreciates over time.[14] The discrete structure of the housing grid gives rise to a minimum house size available for purchase (H_1), meaning that households need to accumulate a certain amount of wealth before they can become homeowners.

---

[12] In the household finance literature, there is an ongoing debate whether stock market participation costs are best approximated by per-period lump-sum costs as in e.g. Vissing-Jorgensen (2002) or Gálvez (2018), or by one-time entry costs, as e.g. in Alan (2006), Cocco (2005) or Gomes and Michaelides (2005). I work with per-period costs to avoid having to introduce risky assets as an additional state variable.

[13] The mortgage premium is constant across all households which is supported by empirical evidence: in Appendix A.1.2, I show that mortgage characteristics of single households do not significantly differ from those of couples in my sample.

[14] For simplicity, I abstract from house price risk. See Appendix A.1.2 for a more detailed discussion.

For homeowners, their house serves as collateral for mortgages. Housing is illiquid, meaning that households have to pay a fraction of the house price whenever they sell or purchase a home. Additionally, they have to pay annual maintenance costs which captures both actual maintenance works but also other housing-related flow expenses such as property taxes. Renting households have to pay a fraction α_R of the price of the smallest owner-occupied house ($p_h H_1$) as rent, with this fraction depending on the specific rental they live in.

1.3.5 Marriage and Divorce

The Evolution of Marital Transitions. Marriage and Divorce are treated as exogenous shocks. Each period, single individuals get married with a probability $\mu(i, j, \tilde{y}_i)$ that depends on their gender i, age j and current productivity realization $\tilde{y}_i$, forming expectations about their prospective partner's asset and income levels.[15] Couples face an age and productivity dependent divorce probability $\lambda(j, \tilde{y}_c)$.[16]

Asset Allocation after Marital Shocks. If two individuals get married, they pool their financial wealth. If neither spouse owns a house at the time of the marital shock, the couple starts married life as renters (and can subsequently jointly re-optimize). If one of the spouses owns a house, the renting spouse moves in with the owning partner. If both spouses are homeowners at the time of marriage, the couple moves into the larger house and sells the smaller one.

If owning couples get divorced, they can either liquidate their house or let one of the spouses keep it, depending on what yields the highest joint continuation value. In the former case, after having liquidated their house, they split all assets equally with a fraction of them being destroyed to account for e.g. legal fees. If one of the spouses keeps the house, the other spouse receives a larger fraction of the couples' financial assets (after an exogenous fraction has been destroyed). All couples who hold negative financial wealth have to liquidate their house. This assumption is necessary to avoid situations in which one spouse receives the entire wealth following a divorce.[17] Renting couples split their financial assets equally upon divorce, again with a fraction of them being destroyed.

15 Section 1.4.1 explains the mapping of partners in terms of observable characteristics in the event of marriage.

16 By targeting marital transitions probabilities conditional on age, income, (and gender), I capture most of the empirical variation in marriage and divorce patterns along observable household characteristics. Nevertheless, to ensure that I am not missing a quantitatively important link between housing and marital risk (e.g. couples may use their house as a commitment device to avoid a divorce), I re-did the analysis using marital transition probabilities of only homeowners. My results are robust to this modification.

17 Because the LTV requirement has to hold each period, households' net worth is always positive.

1.3.6 Taxes

Households pay flat capital taxes τ_k on capital income both from safe and risky assets. Labor income is subject to a progressive tax which maps pre-tax earnings y into post-tax earnings $\mathbb{Y}(y)$ according to[18]:

$$\mathbb{Y}(y) = \tau_l y^{1-\tau_p}$$

The term τ_l governs the average level of taxation and τ_p determines its degree of progressivity. As in the US tax code, mortgage payments above the standard deduction are deductible from the income tax, hence reducing the taxable amount of income, y.

1.3.7 Timing

In the beginning of period t, households learn their current productivity state, their stock market return, and their marital status. Thus, agents start period t with a given amount of net worth that depends on their decisions in period $t-1$, their marital status and the realization of shocks. Afterwards, they decide on how much to consume, their housing stock next period, whether they want to take out a mortgage, and how much to save in risky and safe assets. If they invest part of their endowment in the risky asset (i.e. if $a_{r_{t+1}} > 0$), they have to pay the participation costs S^F in the current period (t).

1.3.8 Recursive Formulation

There are six value functions for singles, couples, and individuals living in couples, both during working age, as well as during retirement.[19] Given that mortgages are modeled as one-period debt, that the stock market participation cost has to be paid per-period, and the i.i.d nature of the return process for the risky asset, I can combine financial assets and labor income into one "cash-on-hand" state variable: $a = \sum_{l=r,s}(1+(1-\tau_k)r_l)a_l - (1+r_m)m + \mathbb{Y}(y(.),m)$ where $\mathbb{Y}(.)$ denote after-tax earnings as described in section 1.3.6.[20]

18 This specification follows Benabou (2002), Heathcote et al. (2017) and Guner et al. (2014).

19 The latter is the relevant object to compute the continuation values of singles in the case of marriage (Borella et al., 2019).

20 Because labor income is not i.i.d, I still keep track of the current productivity realization $\tilde{y}$ when expressing the problem recursively.

Singles – Working Age. The state variables of a single agent are gender i, age j, cash-on-hand a, house $\mathcal{H}$ (which can, in the case that $\mathcal{H} = R$, be rented) and stochastic productivity realization $\tilde{y}$.[21] Each period, she decides on consumption, the housing stock next period, how much to borrow in mortgages, and how much to invest in safe and risky assets. The corresponding value function reads as:

$$V^S(i,j,a,\mathcal{H},\tilde{y}_i) = \max_{a'_r,a'_s,\mathcal{H}',m',c} u(c,s) + \beta(1-\mu(i,j,\tilde{y}_i))\mathbb{E}V^S(i,j+1,a',\mathcal{H}',\tilde{y}'_i)$$
$$+\beta\mu(i,j,\tilde{y}_i)\mathbb{E}\hat{V}^C(j+1,\tilde{a}',\tilde{\mathcal{H}}',\tilde{y}'_c)$$

$$a'_r + a'_s - m' + c = a + p_h\mathcal{H} - p_h\mathcal{H}' - \underbrace{\mathbb{1}_{\mathcal{H}'\neq\mathcal{H}}\Phi(\mathcal{H},\mathcal{H}')}_{\text{Adjustment cost}} - \underbrace{\mathbb{1}_{a'_r>0}S^F}_{\text{SMP cost}} - \underbrace{\mathbb{1}_{\mathcal{H}=R}\alpha_R p_h H_1}_{\text{Rent}} - \underbrace{\mathbb{1}_{\mathcal{H}\neq R}\pi\mathcal{H}}_{\text{Maintenance cost}}$$

$$\underbrace{m' \leq \zeta_h p_h \mathcal{H}'}_{\text{LTV - Constraint}}$$

$$c \geq 0 \qquad a = \underbrace{\sum_{l=r,s}(1+(1-\tau_k)r_l)a_l - (1+r_m)m + \mathbb{Y}(y(i,j,\tilde{y}_i),m)}_{\text{“cash-on-hand”}}$$

where $\tilde{a}'$ and $\tilde{\mathcal{H}}'$ refer to expected financial assets and housing stock, respectively, in the next period if the individual gets married with probability $\mu(i,j,\tilde{y}_i)$. The term p_h denotes the current house price, which is zero for rental properties (i.e. if $\mathcal{H} = R$).

Singles – Retirement. During retirement, singles' state space is characterized by gender i, age j, cash-on-hand a, housing stock $\mathcal{H}$ and the last income realization before retirement ($\hat{y}_i$) which is necessary to compute pension payments. In the terminal period (J), agents have to re-pay all their debt. The term ψ_{ij} denotes age and gender specific survival risk.

$$V^S_R(i,j,\mathcal{H},a,\hat{y}_i) = \max_{a'_s,a'_r,\mathcal{H}',m',c} u(c,s)+$$
$$\beta\psi_{ij}\mathbb{E}V^S_R(i,j+1,\mathcal{H}',a',\hat{y}_i) + \beta(1-\psi_{ij})L\frac{(\xi+a'+\mathcal{H}')^{1-\gamma}}{1-\gamma}$$

$$a'_r + a'_s - m' + c = a + p_h\mathcal{H} - p_h\mathcal{H}' - \mathbb{1}_{\mathcal{H}'\neq\mathcal{H}}\Phi(\mathcal{H},\mathcal{H}') - \mathbb{1}_{a'_r>0}S^F - \mathbb{1}_{\mathcal{H}=R}\alpha_R p_h H_1 - \mathbb{1}_{\mathcal{H}\neq R}\pi\mathcal{H}$$

$$m' \leq \zeta_h p_h \mathcal{H}' \qquad m_J = 0$$

$$c \geq 0 \qquad a = \sum_{l=r,s}(1+(1-\tau_k)r_l)a_l - (1+r_m)m + \mathbb{Y}(pen(\hat{y}),m)$$

[21] The term i denotes family type, i.e. single men, single women or couple. However, when considering only singles, family type and gender are interchangeable.

Couples – Working Age. The state variables of a couple can be summarized by age j, joint cash-on-hand a, joint housing state $\mathcal{H}$ and joint productivity realization $\tilde{y}_c$. The corresponding value function reads as:

$$V^C(j,a,\mathcal{H},\tilde{y}_c) = \max_{a_r',a_s',\mathcal{H}',m',c} u(c,s) +$$

$$\beta(1-\lambda(j,\tilde{y}_c))\mathbb{E}V^C(j+1,a',\mathcal{H}',\tilde{y}_c') + \beta\lambda(j,\tilde{y}_c)\mathbb{E}\sum_{i=f,m} V^S(j+1,\tilde{a}',\tilde{\mathcal{H}}',\tilde{y}_i')$$

$$a_r' + a_s' - m' + c = a + p_h\mathcal{H} - p_h\mathcal{H}' - \underbrace{\mathbb{1}_{\mathcal{H}'\neq\mathcal{H}}\Phi(\mathcal{H},\mathcal{H}')}_{\text{Adjustment cost}} - \underbrace{\mathbb{1}_{a_r'>0}S^F}_{\text{SMP cost}} - \underbrace{\mathbb{1}_{\mathcal{H}=R}\alpha_R p_h H_1}_{\text{Rent}} - \underbrace{\mathbb{1}_{\mathcal{H}\neq R}\pi\mathcal{H}}_{\text{Maintenance cost}}$$

$$\underbrace{m' \leq \zeta_h p_h \mathcal{H}'}_{\text{LTV - Constraint}}$$

$$a = \underbrace{\sum_{l=r,s}(1+(1-\tau_k)r_l)a_l - (1+r_m)m + \mathbb{Y}[y_c(j,\tilde{y}_c),m]}_{\text{“cash-on-hand”}}$$

Again, $\tilde{a}'$ and $\tilde{\mathcal{H}}'$ denote expected financial assets and housing, respectively, in the following period if the couple gets divorced with probability $\lambda(j,\tilde{y}_c)$.

Couples – Retirement. Retired couples individually face the risk of dying. If one spouse dies, the surviving one continues his or her life as single with a fraction δ of the couple's assets and – if they are homeowners – keeps the house. If both spouses die within the same period, they jointly value leaving bequests. Their value function reads as:

$$V_R^C(j,a,\mathcal{H},\hat{y}_c) = \max_{a_s',a_r',\mathcal{H}',m',c} u(c,s) + \beta\psi_{jf}\psi_{jm}\mathbb{E}V_R^C(j+1,a',\mathcal{H}',\hat{y}_c) +$$

$$\beta\sum_{i=f,m}\psi_{ij}(1-\psi_{-ij})\mathbb{E}V_R^S(i,j+1,\delta a',\mathcal{H}',\hat{y}_i) +$$

$$\beta(1-\psi_{jf})(1-\psi_{jm})L\frac{(\xi+a'+\mathcal{H}')^{1-\gamma}}{1-\gamma}$$

$$a_r' + a_s' - m' + c = a + p_h\mathcal{H} - p_h\mathcal{H}' - \mathbb{1}_{\mathcal{H}'\neq\mathcal{H}}\Phi(\mathcal{H},\mathcal{H}') - \mathbb{1}_{a_r'>0}S^F - \mathbb{1}_{\mathcal{H}=R}\alpha p_h H_1 - \mathbb{1}_{\mathcal{H}\neq R}\pi\mathcal{H}$$

$$m' \leq \zeta_h p_h \mathcal{H}' \qquad m_J = 0$$

$$c \geq 0 \qquad a = \sum_{l=r,s}(1+(1-\tau_k)r_l)a_l - (1+r_m)m + \mathbb{Y}(pen(\hat{y}_c),m)$$

Value to an individual of becoming a couple. The value of an individual in a couple is the relevant object when computing the value of single i for getting married to partner p, i.e. the present discounted value of the individual's utility in the event of marriage (Borella et al., 2019). Variables denoted with a $\hat{hat}$ indicate optimal allocations computed with the value function for couples, given the respective state variables. The value of an individual in a retired couple $\hat{V}_R^C$ is defined accordingly.

$$\hat{V}^C(i,j,a,\mathcal{H},\tilde{y}_c) = u(\hat{c},\hat{s}) + (1-\lambda(j,\tilde{y}_c)\beta\mathbb{E}\hat{V}^C(i,j+1,a',\mathcal{H}',\tilde{y}_c') + \lambda(j,\tilde{y}_c)\beta\mathbb{E}V^S(i,j+1,a',\tilde{\mathcal{H}}',\tilde{y}_i')$$

1.4 Calibration

I calibrate the model using a two-step strategy as standard in the literature (e.g. Cagetti, 2003; Gourinchas and Parker, 2002). That is, I first calibrate all parameters that can be identified directly from the data and set some other parameters in line with the literature. Then, I internally calibrate the remaining parameters.

1.4.1 Externally chosen Parameters

I calibrate my model to the years from 1989 until 2017. Table 1.2 summarizes all externally calibrated parameters. The housing grid is defined over five discrete choices: two rentals and three sizes for homeowners, that is $\mathcal{H} = \{R_1, R_2, H_3, H_4, H_5\}$. I set the coefficient of relative risk aversion γ to 1.5 and the housing utility share $(1-\omega)$ to 0.1, both values that are common in the housing literature. I borrow the parameter for the bequest intensity $L = 0.128$ and for the luxuriousness of bequests $\xi = 0.73$ from Cooper and Zhu (2016) who estimate both values in the context of a portfolio choice model with CRRA preferences. The average rent-to-price ratio is 0.1, as estimated in Davis et al. (2008). In particular, I assume the rent for the small rental to be 5% of the smallest (owner-occupied) house price and that of the big rental to be 15% of the smallest (owner-occupied) house price. I follow Cocco (2005) and set the annual maintenance costs to be 1% of the house price. The LTV constraint is set to 0.8 (i.e. households can borrow up to 80% of their house value) and adjustment costs are assumed to be 5% of the house price, both values taken from Paz-Pardo (2020).

Labor Income Profiles. Figure 2.4 plots the empirical life-cycle profiles for average household labor income of single men, single women and couples which inform me about the deterministic component of the labor income process.[22] In per-capita terms, couples' household income is lower than single

[22] Appendix B.3.2 explains in detail how I obtain these profiles.

men's until around age 40. In contrast, single women's labor income is always lower than that of couples and lower than that of single men below age 60.

Figure 1.6 Life-Cycle Income Profiles (Deterministic Component)

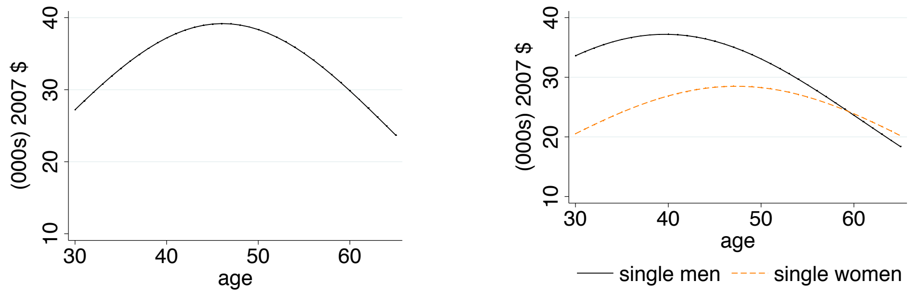

(a) Couples (per capita) **(b)** Singles

Notes: Figure 2.4 plots the life-cycle profiles of the deterministic part of labor income by family type. Labor income is defined as annual earnings out of labor income and social security benefits. Couples' value is expressed in per-capita terms, hence their overall household income is twice as large. Data is from the Panel Study of Income Dynamics (PSID), waves 1989-2017.

The stochastic part of the labor income process displays negative skewness and excess kurtosis across all family types.[23] Both the cross sectional dispersion and the variance in income changes is lower for couples than for singles, suggesting some form of insurance across spouses. For example, couples have the ability to pool individual income streams or to adjust spousal labor supply in response to income shocks, both margins that are not available to singles. In turn, lower income variance affects household's willingness to bear risk along other dimensions, such as asset markets (Fagereng et al., 2018; Heaton and Lucas, 2000). In addition, singles face a higher kurtosis in income changes. Thus, their income process is characterized by more heavy tails, meaning they face larger jumps in their period-by-period income transitions, adding an additional layer of risk.

Pension Payments. Pension payments are assumed to be 70% of the labor income during the last year of work (i.e. at age 64).

Marital Transition Probabilities. I compute marital transition probabilities from PSID data by estimating the following logit function, separately for couples and singles:

$$\xi_{t+1} = \frac{exp(X_t\beta^s)}{1+exp(X_t\beta^s)}$$

where ξ_{t+1} denotes the likelihood of getting married (respectively divorced) within the next period conditional on not being married (respectively being married) in the current period. Explanatory variables include a constant, age, age-squared, current productivity realization and in the case of

[23] In Appendix A.3.2, I report in detail the estimation results as well as the corresponding data fit.

marriage, gender.[24] Figure A.13 and Table A.5 in Appendix B.3.3 report the corresponding life-cycle profiles and regression coefficients, respectively. Both marriage and divorce probabilities are declining in age. In addition, the probability of experiencing a marital transition is non-monotone in income: individuals with medium productivity face the highest marriage probability and are least likely to get divorced whereas individuals at the lowest end of the income distribution are most likely to get divorced and have the smallest probability of getting married.

Marriage Market. Individuals are always matched to a partner with the same age who holds the empirical average amount of financial assets, conditional on age and gender. In 70% of marital unions, the partner is a renter (with a 50:50 chance of living in the small or big rental), whereas the remaining 30% own a small house, corresponding to the average homeownership rate of singles below age 40 (which is when most marriages occur). The probability of meeting a partner such that the couples' productivity realization is $\tilde{y}_c$ depends on the individual's own productivity realization $\tilde{y}_i$ at the time of marriage according to:

$$\Pi_m(.) = \Pi_m(\tilde{y}_c|\tilde{y}_i)$$

I estimate the function $\Pi_m(.)$ non-parametrically from the PSID.

Asset (and Income) Allocation upon Divorce. After a divorce, the first productivity realization as a single depends on the couples' productivity realization at the time of the separation:

$$\Pi_d(.) = \Pi_d(\tilde{y}_i|\tilde{y}_c)$$

which I again estimate non-parametrically from the PSID. Moreover, following Cubeddu and Ríos-Rull (2003), I set the fraction of assets that is exogenously destroyed upon a marital dissolution to 20%. In the event that couples do not liquidate their house, the spouse without the house is left with 70% of the households' financial assets.

Asset Returns. House prices grow deterministically at an annual rate of 3%, which is the average value of the Case-Shiller Index throughout my sample period. The annual return rate of the risk-free asset is 2% and the mortgage premium is 2%, i.e. $r_m = 0.04$. Both values are taken from Cocco (2005). The risky asset has a risk premium of 4%, and a variance of $Var(\tilde{R}(s)) = \sigma_r^2 = (0.1758)^2$, reflecting the annual total return index of the S&P 500 during my sample period. With a 98% probability, the

24 For couples, age refers to the average age across spouses.

return of the risky asset is drawn from that normal distribution and with a 2% probability a disaster state materializes which results in a loss of 40% of all risky assets, both values that Barro (2009) empirically estimates from historical US data on stock market crashes. When simulating the model for a large set of individuals over their life-cycle, I treat the risky asset return as an aggregate shock that evolves according to the observed stock market performance in the US from 1989 until 2016.

Demographic Shifters. Table 1.1 summarizes the values for the demographic shifters η^c and η^s that I obtain from Yang (2009). The first two household members refer to adults whereas all remaining members are children. In the data, I compute the average number of household members by age and family type and assign the corresponding values for both η^c and η^s to each household in the model.

Table 1.1 Equivalence Scales (Yang, 2009)

Family size:	1	2	3	4	5	6	7
η^c (non- housing)	1	1.34	1.65	1.97	2.27	2.57	2.87
η^s (housing)	1	1.1	1.2	1.3	1.4	1.5	1.6

Notes: Table 1.1 lists the demographic shifters for non-durable consumption goods η^c and for shelter services η^s, depending on the number of household members ("family size"). The first two members refer to adults whereas 3 to 7 denote children.

Tax Parameters. I take the values for the tax parameters τ_l and τ_p from Guner et al. (2014) who estimate them using IRS data. I work with their estimates for married couples with one child (the median number of children for couples in my sample), which implies $\tau_l = 0.91$ and $\tau_p = 0.064$, and for singles without children (the median number of children for singles in my sample), resulting in $\tau_l = 0.882$ and $\tau_p = 0.036$.

Survival Probabilities. I compute the gender specific death probabilities at age j from the Life Tables of the US Social Security Administration as the likelihood to die within the next two years conditional on having survived up to age j.[25] I take the inverse of those probabilities and work with average values between the years 1990, 2000 and 2010, corresponding to the sample period of my study. If one member of a couple household dies, the surviving spouse keeps 70% of the household's assets (Jones et al., 2020).

Initial Conditions. The initial distribution of family types mimics the distribution of couples, single men and single women at age 30 from PSID data. The initial distribution of housing is chosen such that it reflects the distribution of homeowners by gender and marital status at age 30 in the SCF. Regarding house sizes, agents initially either rent (with a 50:50 chance of renting the small or the big rental) or

[25] All tables available under this link [Accessed April 19, 2021].

Table 1.2 Externally Calibrated Parameters

Parameter	Source	Value
Model Period Length	PSID frequency	2 years
Housing Grid	–	$\{R_1, R_2, H_3, H_4, H_5\}$
CRRA (γ)	–	1.5
Housing utility share ($1-\omega$)	–	0.1
Bequest Intensity	Cooper and Zhu (2016)	0.128
Luxuriousness of bequest	Cooper and Zhu (2016)	0.73
Rent-to-price ratio (α)	Davis et al. (2008)	0.1
LTV	Paz-Pardo (2020)	0.8
Annual housing maintenance cost	Cocco (2005)	0.01
Housing adjustment cost	Paz-Pardo (2020)	{0.05;0.05}
Survival Probability	Life Tables	see text
Demographic Shifter ($\eta^s < \eta^c$)	Yang (2009)	see text
Tax Parameter	Guner et al. (2014)	see text
Initial Conditions	PSID, SCF	see text
Income Processes	PSID	see text
Prob. of Marriage (μ) & Divorce (λ)	PSID	see text
Asset Returns	Cocco (2005), Barro (2009)	see text

Notes: Table 1.2 lists all model parameters that are either estimated directly from the data or set in line with previous literature.

own the smallest house. The initial distribution of financial assets reflects its empirical counterpart conditional on homeownership status, gender and marital status at age 30 from the SCF.

1.4.2 Internally calibrated Parameters

In the following, I explain the calibration of parameters that cannot be identified directly from the data, with a particular focus on the elasticity of substitution between housing services and non-durable consumption goods (ν).

Elasticity of Substitution between s & c. Large parts the housing literature set the elasticity of substitution between non-durable consumption and housing services to one (i.e. $\nu = 0$) which implies that the momentary utility function $g(c, s)$ takes the Cobb-Douglas form (e.g. Cocco, 2005; Yang, 2009). This assumption can be justified by the almost constant housing expenditure share by wealth and age in micro data (e.g. Davis and Ortalo-Magné, 2011). However, it is no longer suitable for the current set-up because the empirical housing expenditure share of singles is larger than that of couples.[26] Therefore, to pin down ν, I target the ratio of housing expenditure shares between couples and singles. Recall that economies of scale are larger for housing services than for other consumption goods, implying that

[26] See Appendix A.1.2 for more details and corresponding figures.

$\eta^s < \eta^c$. In this case, for the housing expenditure share to be decreasing in the number of household members, it has to hold that $\nu < 0$, meaning that the elasticity of substitution between c and s is below one.[27]

Remaining Parameters. The remaining parameters can be summarized by the discount factor (β), the utility flow from housing services, depending on the specific house size (s_1, s_2, s_3, s_4, s_5), the price of owner-occupied housing (p_3, p_4, p_5) as well as the stock market participation cost (S^F). I normalize the utility flow of the small rental R_1 to one. Hence, including ν, the model has ten free parameters that I jointly calibrate to match ten moments. Table 1.3 summarizes the results.

I target the average net wealth-to-income ratio, that is financial wealth net of mortgages over household income, of couples at age 45 to match the discount factor and take its data value of 1.82 from the SCF. In the model, financial wealth of households is expressed as safe and risky assets net of mortgages which is why the empirical net wealth-to-income ratio is the moment that maps best into the model set-up. I take the homeownership rate of couples at age 45 from the SCF to target the utility flow from living in the larger rental (s_2).

To calibrate the utility flow from owning, I match the average housing share of couples at age 35 (for s_3), at age 45 (for s_4) and at age 55 (for s_5) in the SCF. Importantly, I target both the homeownership rate and housing share of couples because singles' housing tenure choices are more sensitive to the smallest available (owner-occupied) house size (through e.g. lower labor income levels) and hence, I evaluate the model performance by its ability to endogenously replicate their housing choices over the life-cycle (see Section 1.4.3). To pin down house prices, I target average housing wealth (conditional on owning) at different ages. Finally, I match the mean stock market participation rate of couples at age 45 in the SCF to calibrate the flow cost of stock market participation.

Table 1.3 shows that the model matches its associated data targets well. The discount factor ($\beta = 0.888$) is low compared to frameworks with only one financial asset but close to values in the household finance literature with multiple assets. For example, Cooper and Zhu (2016) estimate an annual discount factor of 0.87 in a portfolio framework with CRRA preferences, whereas Catherine (2020) finds $\beta = 0.92$. The estimates for the utility flow of the big rental is $s_2 = 10$. For owner-occupied houses I find $s_3 = 2$, $s_4 = 7$ and $s_3 = 10$. Hence, the per-period flow utility from owing the smallest house is twice as

[27] The optimal relation between non-durable consumption c and housing services s can be expressed as $c = s * \left(\frac{\omega}{1-\omega}\right)^{\frac{1}{1-\nu}} \left(\frac{\eta^s_{ji}}{\eta^c_{ji}}\right)^{\frac{\nu}{1-\nu}}$. Hence, for the housing expenditure share $\left(\frac{s}{s+c}\right)$ to be decreasing in the number of household members, it has to hold that $\nu < 0$ whenever $\eta^s < \eta^c$.

large as renting the smallest rental unit. In contrast, the utility flow from living in the large rental is calibrated to be equal to the utility flow of the biggest house size, allowing households to upgrade their living situation without necessarily having to become homeowners. The calibrated annual stock market participation cost of $1,275 lies within the range of estimates from previous papers that model participation costs as a flow variable, despite a relatively low γ.[28] Cocco (2005) reports an estimate of $1,000 with a coefficient for the relative risk aversion of $\gamma = 5$. Catherine (2020) estimates a stock market participation cost of $1,010 with a CRRA coefficient of $\gamma = 8.2$. In contrast to these papers, the current framework includes marital transition risk, which lowers household's demand for risky assets and thus lowers the calibrated participation cots for a given value of γ. Finally, I find a value for ν of -0.05, implying an elasticity of substitution between non-durable consumption goods and housing services of 0.95.

Table 1.3 Internally Calibrated Parameters: Targets & Fit

Parameter	Value	Key Moment	Data	Model
Discount factor (β)	0.888	mean W/I (net)	1.82	1.83
Big rental size: (s_2)	10	homeownership rate at 45	78%	81%
Small ownership size: (s_3)	2	Housing Share at 35	58%	57%
Medium ownership size: (s_4)	7	Housing Share at 45	61%	53%
Big ownership size:(s_5)	10	Housing Share at 55	55%	42%
Price of small house (p_3)	$120,000	house value of owners at 35	$204,214	$146,552
Price of medium house (p_4)	$180,000	—‖— at 45	$238,085	$184,264
Price of big house (p_5)	$255,000	—‖— at 55	$239,957	$216,708
Stock market cost (S^F)	$1,275 p.a.	mean SMP at 45	62%	62%
Elasticity of subs. (ν)	-0.05	$\frac{\text{hous. expenditure share singles}}{\text{hous. expenditure share couples}}$	1.0860	1.0743

Notes: Table 1.3 lists all model parameters that are internally calibrated to match the moment listed in column "Key Moment". The homeownership rate at age 45, the housing share as well as the mean stock market participation at age 45 refer to couple households.

1.4.3 Model Validation

With the calibrated model at hand, I simulate a panel of 50,000 households over their life-cycle. Using this simulated panel, I validate the model performance by showing its fit for some important untargeted data profiles.

[28] The larger γ, the more risk averse are agents and hence, lower stock market participation costs are needed to match empirical participation rates.

Figure 1.7 Portfolio Allocation around Marital Shocks – Data vs. Model

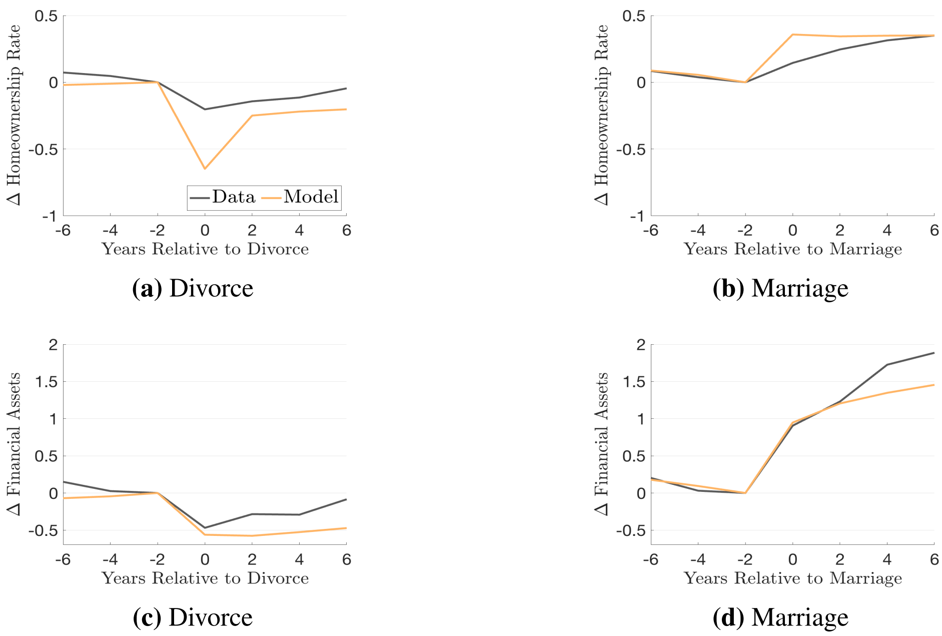

(a) Divorce **(b)** Marriage **(c)** Divorce **(d)** Marriage

Notes: Figure 1.7 plots the change in homeownership rates and in median financial assets in the years preceding and following a marital transition, with values in the year prior to the transition normalized to zero. The gray lines refer to the data (waves 1989 to 2017 of the Panel Study of Income Dynamics (PSID)), whereas the orange lines plot model simulations.

Asset Shifts around Marital Transitions. To validate parameters that govern the marriage market and asset allocations upon marital transitions, Figure 1.7 shows how the model replicates housing choices and changes in financial wealth in the years preceding and following a marital shock, with values in the year prior to the marital transition normalized to zero.[29] Correctly capturing portfolio shifts around the timing of marriage and divorce is crucial to realistically replicate the financial riskiness of marital shocks, which in turn directly affects savings behavior and investment choices of households.

The model captures well the increase in homeownership rates after marriage and the evolution of financial wealth in the event of both marriage and divorce. In contrast, it over-predicts the drop in homeownership rates after a divorce which is partly mechanical: as at most one spouse can keep the house following a divorce, the model naturally produces a drop in homeownership of around 50%pts.[30] Nevertheless, given that it generates an increase in homeownership rates close to empirical levels after one model period (two years), I regard this validation exercise as successful.

Life-Cycle Profiles of Housing and Asset Accumulation. Figure 1.8 shows the model fit for life-cycle profiles of financial wealth accumulation across family types and Figure 1.9 compares homeownership

29 Because the SCF is a repeated cross-section and the PSID has a panel structure, I compute the empirical moments from the PSID despite matching homeownership rates and financial assets from the SCF.

30 The drop would be exactly 50%pts if all divorcees decide that one spouse keeps the house and if the homeownership among couples were 100%.

rates for single men, single women and couples in the data with model-implied simulations. Figure 1.10 reports the model fit for the average housing wealth of homeowners. The model matches very well the financial wealth accumulation of couples and single men whereas it slightly over-predicts the wealth accumulation of single women. Moreover, it is able to replicate the life-cycle path of homeownership rates across all family types and that, conditional on owning, couples live (per capita) in smaller houses. It predicts too little housing wealth for owning couples, thus overstating the (reverse) marital gap in conditional housing wealth. However, most importantly, the model is able to endogenously generate the empirical marital gaps highlighted in Section 1.2: couples are more likely to be homeowners than singles but they live, conditional on owing, in (per capita) smaller houses. In contrast, couples accumulate (per capita) more financial wealth than singles.

Figure 1.8 Financial Wealth by Family Type – Data vs. Model

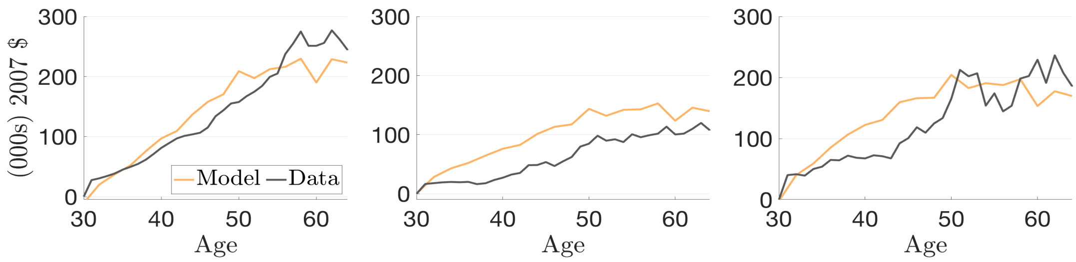

(a) Couples (per capita) **(b)** Single Women **(c)** Single Men

Notes: Figure 1.9 plots the model fit for life-cycle profiles of financial wealth by family type. The gray lines refer to the data (waves 1989 to 2016 of the Survey of Consumer Finances (SCF)), whereas the orange lines plot model simulations.

Figure 1.9 Homeownership Rates by Family Type – Data vs. Model

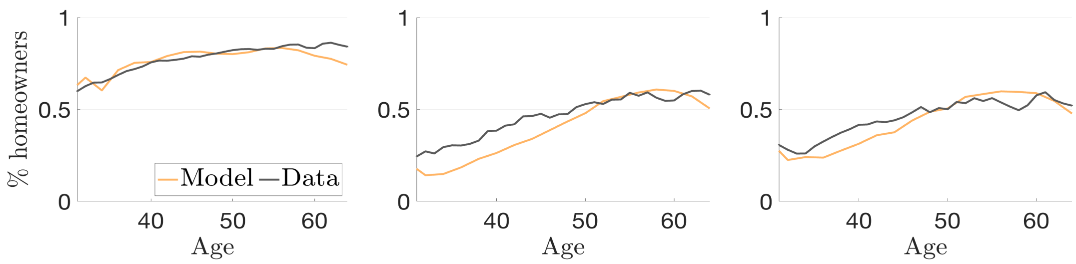

(a) Couples **(b)** Single Women **(c)** Single Men

Notes: Figure 1.9 plots the model fit for life-cycle profiles of homeownership rates by family type. The gray lines refer to the data (waves 1989 to 2016 of the Survey of Consumer Finances (SCF)), whereas the orange lines plot model simulations.

Further Results. Appendix A.4 reports the model fit for the share of overall wealth that households allocate to housing, safe and risky assets by age groups (corresponding to Figure 1.3). In Appendix A.2, I test the sensitivity of my results to increasing the housing grid and validate the model performance with regard to matching empirical moving frequencies by marital status. Previous literature has documented that couples move less often than singles (e.g. Blackburn, 2010; Burke and Miller, 2018; Gemici, 2011;

Mincer, 1978) which could shift their incentive to become homeowners. To further address this concern, I conduct a robustness in which I introduce an iid moving shock as in Cocco (2005) that is allowed to differ by marital states. I show that the main results are not sensitive to its introduction.

Figure 1.10 Average Housing Wealth of Owners by Family Type – Data vs. Model

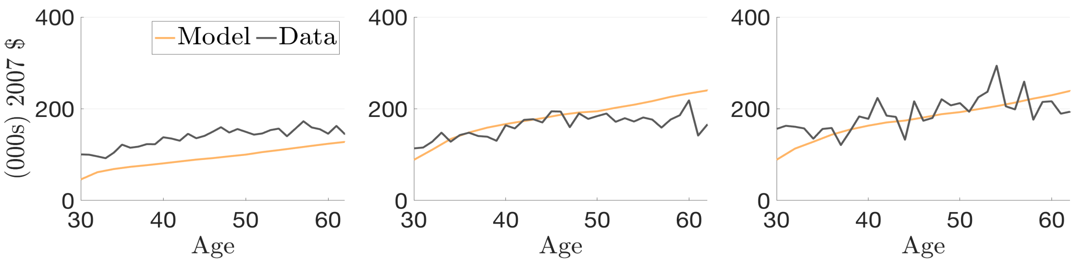

(a) Couples (per capita) **(b)** Single Women **(c)** Single Men

Notes: Figure 1.10 plots the model fit for the average house value of home owners by family type. The gray lines refer to the data (waves 1989 to 2016 of the Survey of Consumer Finances (SCF)), whereas the orange lines plot model simulations.

1.5 (How) Does Marital Status Affect Housing Demand?

By means of counterfactual simulations, I now turn the of the first research question and study the channels through which marital status affects households' investment choices. In each counterfactual, I change one element, re-solve and re-simulate the model and contrast the resulting life-cycle profiles to the baseline economy. To analyze the role of marital risk, I shut down marriage and divorce ($\mu = \lambda = 0$). To further disentangle the relative importance of each factor, I perform one counterfactual with only marriage ($\lambda = 0$) and one with only divorce ($\mu = 0$). I then evaluate the relative contribution of marital heterogeneity in labor income levels, labor income risk and in household sizes (through economies of scale) by changing the value of both single men and single women for each element to the corresponding value of couples.

1.5.1 The Role of Marital Transition Risk

Marriage and divorce are key drivers of the observed marital gaps in investment choices. Figure 1.11 shows the aggregate change in financial wealth accumulation, in homeownership rates and in conditional housing wealth of couples and singles in response to shutting down marital transitions. All changes are expressed in percent. In addition, Figures 1.12 and 1.13 compare the housing policy functions for couples and single men between the baseline model and each marital counterfactual. Appendix A.4.2 reports the corresponding policy functions for single women.

Figure 1.11 Counterfactuals – The Role of Marital Transition Risk

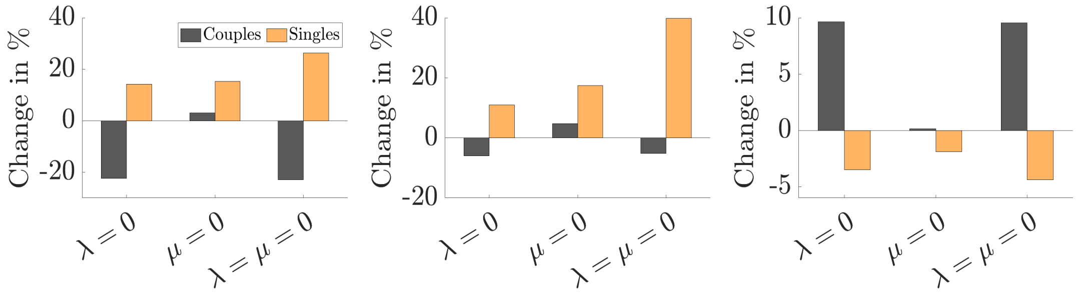

(a) Asset Accumulation **(b)** Homeownership Rate **(c)** (Cond.) Housing Wealth

Notes: Figure 1.11 reports the change in asset accumulation, homeownership rates and conditional housing wealth when shutting down divorce ($\lambda = 0$), shutting down marriage ($\mu = 0$) or both ($\mu = \lambda = 0$). The gray bars refer to couples whereas the orange bars denote singles. All changes are expressed in percent.

Figure 1.12 Housing Policy Functions – Couples

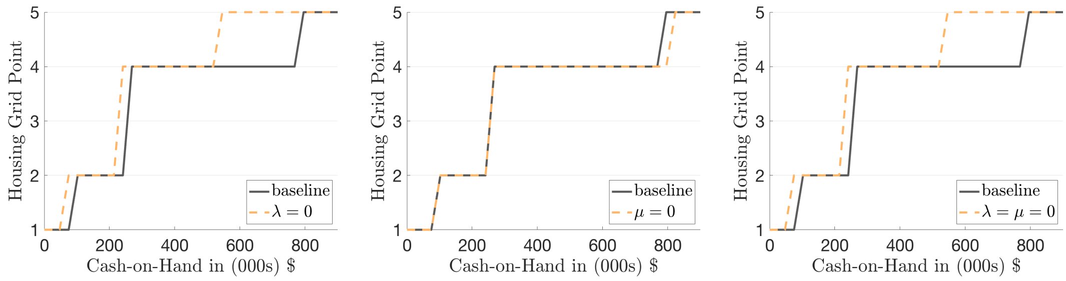

(a) No Divorce **(b)** No Marriage **(c)** No Marital Transitions

Notes: Figure 1.13 plots the housing policy functions for couples in the baseline as well as in the counterfactual without divorce ($\lambda = 0$), without marriage ($\mu = 0$) and without any marital transitions ($\lambda = \mu = 0$). All Figures refer to couples of age 30 who rent the smallest house size and have a medium productivity realization.

No Divorce ($\lambda = 0$). In the absence of divorce risk, couples reduce their (precautionary) financial savings by around 20% (Figure 1.11a). This effect arises from two sources. First, divorce results in a destruction of part of the household's assets against which it wishes to self-insure. Second, once being divorced, couples' exposure to labor income risk increases and, through economies of scale, they need more than half of the previous consumption level to maintain the same level of utility.

Additionally, Figure 1.12a shows that shutting down divorce lowers the asset threshold at which couples transition into ownership and at which they increase their housing size, reflecting an increase in their housing demand. With regard to housing tenure choices, the reduced savings motive is quantitatively stronger than the increased housing demand: as the distribution of couples shifts towards lower-asset households, the aggregate homeownership rate of couples drops by around 7% (Figure 1.11b).

In contrast, as displayed in Figure 1.11c, the conditional housing wealth of couples increases. Because of increased housing demand, equally rich couples now invest in larger houses (Figure 1.12a). Moreover,

Figure 1.13 Housing Policy Functions – Single Men

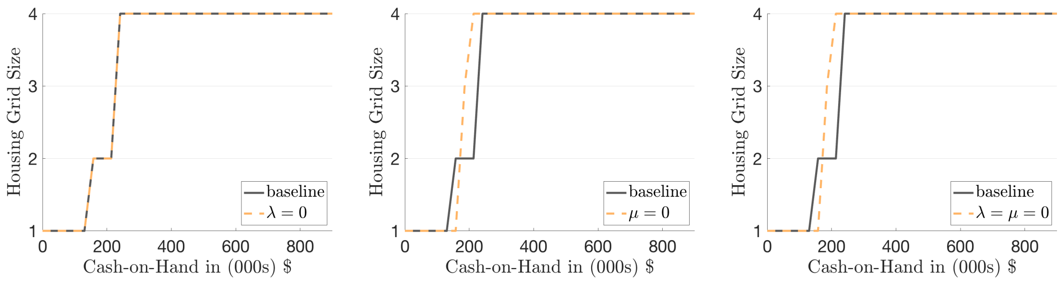

(a) No Divorce **(b)** No Marriage **(c)** No Marital Transitions

Notes: Figure 1.13 plots the housing policy functions for single men in the baseline as well as in the counterfactual without divorce ($\lambda = 0$), without marriage ($\mu = 0$) and without any marital transitions ($\lambda = \mu = 0$). All Figures refer to single men of age 30 who rent the smallest house size and have a medium productivity realization. Appendix A.4.2 reports the corresponding policy functions for single women.

as the homeownership rate of couples drops, some low asset households become renters, shifting the distribution of owning couples towards larger homes.

Aggregate financial savings of singles increase by 15% which arises from a composition effect rather than from a change in individual housing demand (Figure 1.13a): never married singles hold on average more financial assets than divorced individuals because low income households are more likely to divorce and because divorce is costly.[31] Consequently, the share of single homeowners increases and the distribution of owning singles shifts towards smaller houses, resulting in lower conditional house values.

No Marriage ($\mu = 0$). For singles, marriage acts as a financial outcome with disproportionally high returns through asset holdings of the prospective partner and allows for the possibility to pool income as well as to benefit from economies of scale. As a result, singles accumulate more financial assets in a world without marriage than they do in the baseline (Figure 1.11a). In addition, Figure 1.13b shows that, conditional on their cash-on-hand level, singles are more likely to become homeowners as they no longer face the possibility of meeting a partner and having to sell the house. Quantitatively, the homeownership rate of singles increases by around 17% (Figure 1.11b). In contrast, as some previously renting singles now own the smallest owner-occupied house, the conditional housing wealth of singles declines (Figure 1.11c).

The stronger savings motives of singles induces couples to save more as well because they want to hold sufficient financial assets in the event of divorce (Figure 1.11a). As a result, the homeownership rate

[31] The fact that never married singles hold more financial wealth than divorced individuals is in line with empirical evidence, see Figure A.1 for more details.

of couples increases slightly. However, as these changes are quite small, the house value of owning couples remains almost the same as in the baseline (Figure 1.11c).

No Marital Transitions **(**$\mu = \lambda = 0$**).** For couples, the effect of shutting down divorce is quantitatively so much stronger than shutting down marriage that their response in the counterfactual without any marital transitions remains virtually the same as in the one with only marriage, both on aggregate as well in terms of policy functions. For singles, by contrast, financial savings increase on aggregate by 26%, compared to 15% in each of the previous two counterfactuals. This result reflects a combination of their increased savings motive in the absence of marriage and the fact that never married singles hold are on average more financial assets than divorced individuals. In addition, the aggregate homeownership rate of singles increases substantially (by almost 40%). Finally, the average housing wealth of owning singles decreases because many low-asset households switch from renting to owning the smallest house, tilting the conditional distribution of housing wealth towards smaller homes.

Figure 1.14 Counterfactuals – Further Channels

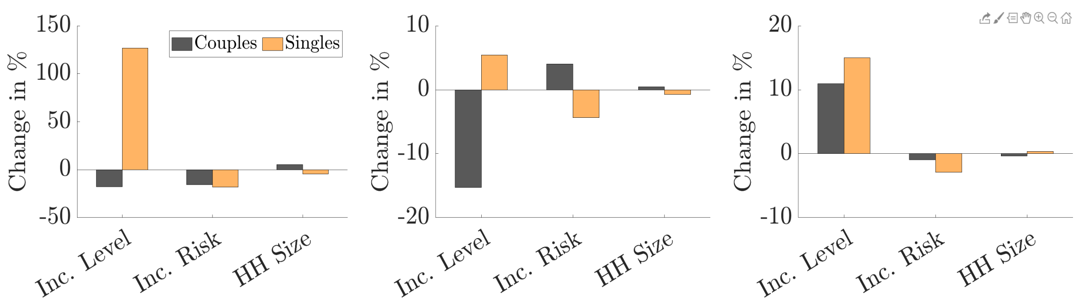

(a) Asset Accumulation **(b)** Homeownership Rate **(c)** (Cond.) Housing Wealth

Notes: Figure 1.14 plots the change in asset accumulation, homeownership rates and conditional housing wealth when assigning singles the deterministic part of couples' income process ('Inc. Level'), the stochastic part of couples' income process ('Inc. Risk') and their average household sizes, conditional on age ('HH size'). The gray bars refer to couples whereas the orange bars denote singles. All changes are expressed in percent.

1.5.2 Further Factors

To explore other channels that further contribute to marital gaps in investment choices, Figure 1.14 plots the change in financial wealth accumulation, in homeownership rates and in conditional housing wealth in response to changing singles' labor income profiles (separately for the deterministic and the stochastic component) and their average number of household members to the corresponding couple value. Again, all changes are expressed in percent.

Income Level. Assigning singles the deterministic part of couples' income process effectively increases their average labor income. Consequently, singles save more, are more likely to be homeowners and live in larger houses. Couples, in contrast, save less than in the baseline. The income drop in the event of divorce becomes smaller and hence, divorce is not as financially risky. Furthermore, through reduced (marital) risk exposure, the housing demand of couples increases.[32] With regard to housing tenure choices, the reduced savings motive again dominates the increased housing demand, resulting in a lower homeownership rate of couples (Figure 1.14b).

Income risk. When assigning singles the stochastic part of couples' labor income process, I lower their exposure to income fluctuations. As a result, precautionary savings of singles decline (Figure 1.14a). In addition, because of reduced labor income risk, the housing demand of singles increases. Nevertheless, through fewer financial savings, homeownership rates of singles slightly drop on aggregate and the distribution of single owners shifts to slightly smaller houses. For couples, divorce is again a less financially risky outcome as in the baseline. Consequently, they accumulate less financial assets and their demand for housing increases. In contrast to singles, the share of owning couples becomes larger, as the larger willingness to invest into housing dominates the reduced savings motive.

Household sizes. When assigning singles the same average household members as couples, I increase their household size, conditional on age. In response, singles have larger consumption needs each period, resulting in lower financial savings, lower homeownership rates and in a slight increase of conditional house values. For couples, divorce becomes more risky, and in response, their financial savings and homeownership rates increase. However, overall, the effect of changing household sizes is small, especially when compared to the importance of marital transition risk.

1.6 Implications for Policy Evaluation

In this section, I address the second research question and show that abstracting from distinct family types is misleading in judging the effectiveness of policies that aim at stimulating homeownership, especially among young households, which is the age group that most housing reforms in the US are primarily targeted to.

First, I lower the transaction costs of housing Φ from 5% to 2% of the house price. Second, I reduce property taxes by decreasing annual maintenance costs π from 1% to 0.45% of the house price.[33] Hence, the first policy change facilitates housing adjustment in response to shocks whereas the latter

32 The corresponding policy functions are reported in Appendix A.4.2.
33 In Appendix A.5.3, I show that the results in this section are robust to changes in house prices and to changes in marital transitions probabilities in response to the introduction of both reforms.

aims at lowering the flow cost of housing. To make both policies comparable in magnitude, I require the average per-household gain to be similar across reforms. For example, when lowering housing adjustment costs from 5% to 2% of the house price, I calculate the overall "savings" on all housing transactions that occur in the economy after the policy implementation (i.e. 3% of the respective house price per transaction) and average these savings across all years and households.

To analyze the importance of marital risk and family composition, I perform the same policy exercise in a standard framework with one generic household type. To do so, I collapse all three family types and re-calibrate the income process, household sizes (i.e. the demographic shifters), tax parameters and survival risk for the pooled sample while fixing preference parameters to be the same as in the benchmark.[34]

1.6.1 Increasing Homeownership

Table 1.4 displays the increase in homeownership rates across family types in response to both policies. The row "Annual per-HH Gain" reports the described measure of magnitude. *Panel I* shows the results for the benchmark economy whereas *Panel II* displays the results for an economy with distinct family types but without marital risk, i.e. the benchmark framework with $\mu = \lambda = 0$. *Panel III* contains the results for the reduced framework.

In the benchmark economy, both policies result in a quantitatively similar increase of homeownership rates by around 5.5%pts and this increase is evenly distributed across household types: for couples, both policies lift homeownership rates by around 6%pts, whereas the share of single homeowners increases by 4-5%pts.

In contrast, when shutting down marital transitions (*Panel II*), homeownership rates increase by 8.46%pts when lowering adjustment costs and by 11.91%pts when lowering property taxes. In the absence of marriage and divorce, households are more willing to invest their wealth into (illiquid) housing. Consequently, both policies attract more home-buyers who previously preferred to remain renters.

Moreover, lowering property taxes ($\pi \downarrow$) now appears to be over 40% more effective than facilitating house size adjustments ($\Phi \downarrow$). Lower maintenance costs decrease the per-period expenditure commitments of homeowners and thus, make them less vulnerable to income fluctuations. However, once I

34 The reduced framework matches key aggregate data moments of asset accumulation and housing choices when using the same parameters as in the benchmark. See Appendix A.6 for details.

Table 1.4 Effectiveness of Housing Policies Across Frameworks

	Δ Homeownership Rate	
	$\Phi \downarrow$ $(5\% \rightarrow 2\%)$	$\pi \downarrow$ $(1\% \rightarrow 0.45\%)$
Annual per-HH Gain:	\$400	\$390
Panel I: Benchmark		
Couples	+6.03%pts	+5.88%pts
Single Men	+4.30%pts	+4.73%pts
Single Women	+3.89%pts	+5.21%pts
Aggregate	+5.52%pts	+5.64%pts
Panel II: $\lambda = \mu = 0$		
Couples	+9.16%pts	+13.02%pts
Single Men	+4.50%pts	+4.23%pts
Single Women	+4.11%pts	+6.64%pts
Aggregate	+8.46%pts	+11.91%pts
Panel III: One HH-Type	+8.43%pts	+13.16%pts

Notes: Table 1.4 reports the average increase in homeownership after lowering housing transaction costs ($\Phi \downarrow$) and lowering housing maintenance costs ($\pi \downarrow$) in the benchmark *(Panel I)*, in the benchmark without marriage and divorce *(Panel II)* and in the reduced framework with one generic household type *(Panel III)*.

account for marital transitions, households face the risk of having to sell their house (either following a divorce or because they move in with their partner), increasing their desire for being able to do so at little cost.

Furthermore, the increase in households' responsiveness across *Panel I* and *Panel II* is almost entirely driven by couple households. Singles have lower income levels than couples and are exposed to more labor income fluctuations. Thus, even in the absence of marriage, the share of singles who either cannot afford or do not want to buy a house (to be better able to smooth consumption in response to income shocks) remains relatively large. In contrast, most couples have the financial means to invest in owner-occupied housing and it is rather the possibility of divorce (which requires allocating or liquidating their home) that makes them reluctant to become owners.

As a result, the reduced framework which abstracts from both marital risk and distinct family types (*Panel III*) predicts an even stronger increase in homeownership, especially in response to the second reform. Compared to the benchmark, it overstates the effectiveness of lowering transaction costs by 53% and of lowering maintenance costs by 133%.

1.6.2 Heterogeneity over the Life-Cycle

In this section, I explore whether the magnitude of the policy mis-specification between the benchmark and the reduced framework varies over the life-cycle. Table 1.5 compares the increase in homeownership rates in response to both reforms across set-ups for young (age 30 to 39), middle-aged (age 40 to 49) and old (age 50 to 64) households.

In both economies, the effect of the policies become stronger as households age, suggesting that the overall increase in the share of homeowners is not merely driven by earlier transition into ownership. Additionally, the discrepancy across frameworks is strongest early in the life-cycle: with regard to lowering transaction costs, abstracting from family types overstates the policy response of households below age 40 by 108%, of middle-aged households by 76% and of old households by 41%. The intuition behind this result is twofold. First, marriage and divorce probabilities are declining in age. Hence, abstracting from marital risk increases the attractiveness of housing investments the most for young households who consequently react more strongly to the introduction of the policy. Second, the share of single households – who are least responsive to housing policies – is largest among the age group below 40, further contributing to the negative age gradient of the mis-specification. However, in the US, most housing policies are primarily targeted towards young households, further emphasizing the importance of taking into account family composition when evaluating such reforms.

1.6.3 Fostering Overall Wealth Accumulation

Enabling more households to become homeowners is often regarded as desirable because housing represents an important channel of wealth accumulation for middle-class Americans. Therefore, I now turn to evaluating the proposed policy reforms in terms of increasing households' net worth. Additionally, I study the effect of lowering stock market participation costs. Table 1.6 reports the results.[35]

In the benchmark framework (*Panel I*), lowering stock market participation costs is most effective in terms of fostering overall wealth accumulation and increases average household net worth by $8,737. In contrast, both housing policies do so only by a little over $5,000. This effect is especially pronounced for single households: encouraging stock market participation increases the average net worth of single men by 228% more than fostering housing investment and the average net worth of single women by

[35] Table A.8 in Appendix A.7 splits the increase in overall net worth into changes in average house values and changes in aggregate financial savings.

Table 1.5 Effectiveness of Housing Policies by Age Groups

	Δ Homeownership Rate		
Age	30 to 39	40 to 49	50 to 64
Panel I – Housing transaction costs $\downarrow$:			
Couples	+2.44%pts	+4.70%pts	+8.88%pts
Single Men	+2.98%pts	+4.06%pts	+5.57%pts
Single Women	+0.78%pts	+4.84%pts	+6.18%pts
Aggregate	+2.34%pts	+4.64%pts	+7.95%pts
One HH-Type	+4.86%pts	+8.17%pts	+11.21%pts
Difference	108%	76%	41%
Panel II – Housing maintenance costs $\downarrow$:			
Couples	+2.56%pts	+4.26%pts	+8.75%pts
Single Men	+3.10%pts	+5.17%pts	+5.79%pts
Single Women	+2.63%pts	+6.56%pts	+6.91%pts
Aggregate	+2.61%pts	+4.67%pts	+8.01%pts
One HH-Type	+9.53%pts	+16.95%pts	+12.51%pts
Difference	265%	263%	56%

Notes: Table 1.5 reports the average increase in homeownership after lowering housing transaction costs ($\Phi \downarrow$, *Panel I*) and after lowering housing maintenance costs ($\pi \downarrow$, *Panel II*) in the benchmark and in the reduced framework with one generic household across different age groups. The columns "Difference" display the increase in homeownership rates in the reduced framework when compared to the aggregate increase in the benchmark.

123% more, compared to a 31% increase for couples. Singles have lower labor income than couples, keeping them out of the stock market (due to participation costs) and out of homeownership (due to a minimum house size). Reducing stock market participation costs enables them to enter the stock market and to invest relatively little wealth into risky assets. In contrast, even with reduced transaction or maintenance costs, becoming a homeowner still requires relatively large amounts of wealth to pay for the downpayment.

When turning to the reduced framework with one generic household-type, *Panel II* in Table 1.6 shows that lowering maintenance costs increases average household net worth by a little more than $7,000, decreasing transaction costs by around $3,000 and facilitating stock market participation by $5,769. Hence, encouraging investment in risky financial assets does not necessarily appear to be more effective in terms of fostering overall wealth accumulation, again altering the results drawn from the benchmark framework with distinct family types.

Table 1.6 Effect of Housing Policies on Net Worth

	Δ Net Worth in \$		
	$\Phi \downarrow$ $(5\% \to 2\%)$	$\pi \downarrow$ $(1\% \to 0.35\%)$	$S^F \downarrow$ $(\$1,275 \to \$713)$
Annual per-HH Gain:	$400	$390	$395
Panel I: Bench			
Couples	8,427	6,868	10,041
Single Men	-3,701	-1,571	5,996
Single Women	1,347	4,175	6,161
Aggregate	5,316	5,097	8,737
Panel II: One HH-Type	2,945	7,015	5,769

Notes: Table 1.6 reports the average increase in households' net worth in response to lowering housing transaction costs ($\Phi \downarrow$), lowering housing maintenance costs ($\pi \downarrow$) and lowering stock market participation costs ($S^F \downarrow$) in the benchmark economy *(Panel I)* and in the reduced framework *(Panel II)* with one generic household type.

1.7 Conclusion

This paper analyzes how marital status interacts with housing decisions of individuals and shows that explicitly taking family structure into account is necessary for the correct evaluation of policies that aim at stimulating housing demand, especially early in the life-cycle.

First, I provide novel empirical evidence that singles are less likely to be homeowners than couples but that they – conditional on owning – allocate more wealth into housing. In contrast, couples accumulate per capita more financial wealth than singles. By developing a life-cycle framework of family types, housing and financial portfolio choice, I show that low income levels of singles and the presence of marriage and divorce induce couples to accumulate more (precautionary) savings whereas it depresses savings of singles, contributing to the marital gap in financial wealth and in homeownership rates. Lower income risk of couples decreases the asset threshold at which they become homeowners, shifting the distribution of owning couples towards smaller houses. Abstracting from distinct family types biases the effectiveness of policies intended to increase homeownership as it overstates the attractiveness of illiquid housing. This bias is most strongly pronounced among young households whose marital transition risk is highest and among whom the share of single households is largest. However, they are the primary target group of housing policies in the US, highlighting the importance of taking into account marital status when evaluating or designing such reforms.

Chapter 2

The Gender Investment Gap over the Life-Cycle

Abstract Single women are less likely to hold risky assets than single men and allocate a smaller share of their portfolio into stocks. This paper develops and estimates a portfolio choice model to quantify the determinants of the "gender investment gap" over the life-cycle. The framework allows for differences in household structure (single or couple), marital transitions as well as for rich gender heterogeneity along observable characteristics and stochastic processes. The model is able to rationalize the gender gap in equity shares and in asset holdings without introducing preference heterogeneity across men and women. Counterfactual simulations reveal that both current and expected lower income levels as well as larger household sizes of single women are the main determinants for explaining the investment gap. In particular, expectations about future income levels and household sizes drive most of the investment differences for young individuals whereas heterogeneity in current income levels (and household sizes) explain the gender investment gap later in life.

2.1 Introduction

Single women are less likely to participate in the stock market than single men and if they do, they allocate a smaller share of their portfolio towards risky assets. However, in the presence of an equity premium and diversification gains, a more conservative portfolio translates (ceteris paribus) into lower wealth levels. This paper studies the sources of the so-called "gender investment gap" based on an estimated structural life-cycle framework. Generally, differences in investment behavior can arise due to differences in circumstances (such as income levels, income risk, household size etc.) or due to differences in unobservable characteristics such as preferences. In this paper, I ask how much of the gender investment gap can be explained by the former.

To that end, I first document life-cycle profiles of asset holdings, equity shares, stock market participation rates and equity shares conditional on participation ("conditional risky shares") for single men, single women and for couples using survey data on US households who were born between 1945 and 1960. My empirical findings confirm the gender investment gap: Women are less likely to participate in the stock market and allocate – conditional on participating – a lower share of their portfolio towards risky assets. All differences are statistically different from zero, even after controlling for a wide range of observables such as age, education, labor income and the number of household members. In particular, the unexplained part of the gender investment gap decreases as households age.

To uncover which factors explain the remaining gap and to quantify the relative importance of each channel, I go on to develop a life-cycle model of portfolio choice that allows for marital transitions over the life-cycle, for differences in household structure (single or couple) and in gender. In the model, single men and single women differ with regard to their income levels, their income risk, the number of individuals who live in their household, marital transitions probabilities, the (expected) characteristics of their partner in the event of marriage, their survival probabilities as well as their out-of-pocket medical expenditures during retirement. I restrict preferences to be homogeneous across single men and single women to study what fraction of the gender investment gap can be explained by differences in circumstances within the structural model. In contrast, I do allow for preference heterogeneity by marital status (i.e. between couples and singles) in order to better accommodate the data while at the same time keeping the model tractable.

Next, I estimate the model using data from the Survey of Consumer Finances (SCF) for financial choices and from the Panel Study of Income Dynamics (PSID) for labor income and demographic characteristics. To do so, I first estimate all parameters that can be cleanly identified outside of the model and afterwards estimate the remaining parameters using the Simulated Method of Moments (SMM), taking first-stage parameters as given. The model matches well the life-cycle profiles of wealth holdings and equity shares for single men, single women and couples. Finally, I decompose the gender investment gap along the dimensions of gender heterogeneity within the model by replacing the female values with that of single men.

The main results are as follows. First, heterogeneity in income levels accounts for almost 40% of the gender gap in equity shares. Thus, the gender wage gap is amplified through a less risky investment strategy, paying on average lower returns. Thereby, not only the current period income matters but also the fact that single men are endowed with more human capital (i.e. they expect higher income in future periods). Cocco et al. (2005) show that labor income risk is uncorrelated to asset returns and therefore, it acts as a substitute for the safe asset. Consequently, a higher human capital endowment increases the willingness to take on financial risk for any given level of current labor income (and other state variables, such as wealth). Moreover, I find that differences in the number of household members are key in explaining the gender gap in equity shares. Over the course of their working life, larger female household sizes – which arises mainly through a higher likelihood of having children – can explain 43.16% of the observed gap. Again, not only the current household size affects savings and equity shares (through different consumption needs) but larger expected household sizes act as a future consumption commitment, making single women more vulnerable to financial shocks and hence, reduce financial risk-taking.

Lastly, I decompose the counterfactual scenario in which I assign single women the male income level into a *composition* effect, that is how much of the differences in female equity shares between the baseline and the counterfactual can be explained by changes in the distribution of individuals across the state space, and into a *policy* effect, that is how much can be explained by differences in the policy functions for equity shares conditional on state variables. My findings suggest that in a world where single women had the same income level as single men, most of the increase in mean female equity shares early in life occurs because of the *policy effect*, that is conditional on state variables (through higher expected income in future periods). At around age 56, this relation flips and most of the increase in equity shares can be attributed to the distribution of individuals across the state space (as higher labor income in past years has translated into larger wealth levels). Hence, reduced form regressions that do not control for expectations about future income have less prediction power in explaining the gender investment gap early in life when these expectations are more important, which is in line with my empirical findings.

Related Literature. This paper contributes to several strands of the literature. First, it adds to a literature documenting gender differences in investment behavior and in financial choices. In general, there is large consensus that women invest less risky than men. For example, Sunden and Surette (1998) and Agnew et al. (2003) show that women in the US choose lower equity allocations in retirement saving plans than men. Similarly, Barber and Odean (2001) find that single men trade more frequently in risky assets than single women and attribute this result to male overconfidence. Säve-Söderbergh (2012) explores how men and women choose risk profiles in their pension contribution plans in Sweden. She documents that even though women do not less frequently include stocks in their portfolio, they do allocate a smaller share into risky assets. In more recent work, Almenberg and Dreber (2015) or Thörnqvist and Olafsson (2019) show that the gender investment gap prevails until today in Sweden. Ke (2018) attributes cross-country differences in stock market participation rates to gender norms, showing that countries with strong gender norms exhibit lower stock market participation rates of women. Moreover, Goldsmith-Pinkham and Shue (2020) provide evidence that women not only invest more conservatively in liquid financial assets but that they also earn lower returns on housing investments. Similarly, Andersen et al. (2020) find gender differences in negotiation outcomes using transaction data of Danish residential real estate. My paper adds to this literature by being the first – to be best of my knowledge – to analyze the gender investment gap in a structural framework. Most of the previous literature focuses on measurement whereas my setting allows to model different channels and to quantify their relative importance. Relatedly, there exists an experimental literature that documents a higher risk aversion for women, also with regard to financial choices (see e.g. Croson and Gneezy (2009) for a review). Even though my model replicates differences in equity shares with homogeneous preferences across gender, my results are not in contrast to previous experiments: Single women behave *observationally* different to single men in the model, conditional on observable characteristics. However, my structural frameworks predicts that it is not underlying preference parameters that drive these differences but rather expectations about lower income levels and larger household sizes (which act as future consumption commitments) in future periods.

Second, my paper relates to a literature that explores how family related shocks (such as marriage or divorce) affect portfolio allocation and savings. Cubeddu and Ríos-Rull (2003) study the role of marriage and divorce on wealth accumulation in a dynamic setting. Love (2010) was the first to present a joint life-cycle framework of marital status and portfolio choice. He finds that married investors hold more risky portfolios than singles. In the event of divorce, stock holdings increase for men whereas they decline for women. Hubener et al. (2015) extend the analysis by incorporating endogenous labor supply and realistically calibrated social security benefit claiming. Again, they show that the equity share of couples is higher than for singles and that uncertain fertility can significantly reduce the amount of stock holdings. Christiansen et al. (2015) empirically address the heterogeneous impact of family shocks on portfolio choices across gender with a difference-in-difference approach using an administrative panel dataset from Denmark. Similar to Love (2010) for the US, their findings suggest that the fraction of risky assets in women's portfolio increases after marriage whereas it declines after a divorce. For men, this relationship points in the opposite direction. Along the same lines, Bertocchi et al. (2011) find in an empirical framework that the marital gap of stock holdings (i.e. that stock holdings tend to be higher for married than for single individuals) in Italy is larger for women than it is for men. However, while all these papers show that family-related shocks affect portfolio choices heterogeneously across gender, neither of them quantifies the importance of such for the gender investment gap over the life-cycle. Finally, Bogan and Fernandez (2017) find that having a child with mental disabilities decreases stock market participation but increases the share of wealth allocated to risky assets, conditional on participating.

More broadly, my paper extends a literature that studies life-cycle pattern of household finances. Typically, life-cycle models of portfolio choice predict the optimal equity share to be decreasing in the ratio of current financial wealth over the present value of human capital (Merton (1971), Viceira (2001)).[1] Consequently, it should be optimal for young investors (who are endowed with relatively little financial wealth compared to their human capital stock) to allocate 100% of their financial wealth into stocks and to decrease the equity share as they age. In contrast, we observe only limited stock market participation and (conditional) equity shares, especially for young investors, in the data. The literature has proposed several mechanisms to explain this discrepancy. The most prominent explanation are costs associated with stock market investment (Vissing-Jorgensen (2002), Gomes and Michaelides (2005), Alan (2006)). Moreover, papers have highlighted the importance of the illiquid nature of housing (Cocco, 2005), lack of financial literacy (Lusardi and Mitchell, 2014) and cyclicality of labor income (Catherine, 2019).

Finally, my paper is methodologically related to Cooper and Zhu (2016) as they too estimate a life-cycle model of portfolio choice to study why certain subgroups of the population display different investment patterns. In contrast to the present study, the focus of the paper is on education and not on gender. Their findings suggest that income levels are the major determinant why more educated households invest more heavily in risky assets. Similarly, I find that income differences across single men and single women explain the largest part of the gender investment gap. With regard to my modeling approach, this paper is one of the first to introduce couples, single men and single women into a unified portfolio choice framework. I follow Fagereng et al. (2017) in their way to introduce a portfolio choice and build on Borella et al. (2019) to introduce exogenous marital transitions and the assortativeness of couple formation.

[1] The intuition behind this finding is that the correlation between labor income (human capital) and asset returns is almost zero in the data and therefore, human capital acts as a substitute for safe assets (Cocco, 2005).

The remainder of the paper is structured as follows. Section 2.2 presents empirical observations on gender specific portfolio choices. Section 2.3 introduces the structural model. Section 2.4 introduces the estimation strategy and Section 2.5 presents the quantitative results. In Section 2.6, I decompose the gender investment along different sources of gender heterogeneity in the model. Finally, Section 2.7 concludes.

2.2 The Gender Investment Gap in the Data

The following section first explains the data and the sample selection criteria. In a next step, I provide empirical evidence on portfolio choices of single men, single women and couples over their life-cycle.

2.2.1 The Sample

Throughout the analysis, I restrict the sample to individuals from 30 to 65 years who were born between 1945 and 1960. Ameriks and Zeldes (2004) point out that pooling multiple cohorts results in different life-cycle profiles for investment choices depending on whether one controls for time or for cohort effects. Therefore, I focus on individuals born within a relatively short time frame while controlling for age and time effects. In that way, I ensure that all individuals in my sample faced similar environmental conditions at a given age.

I use the waves 1989 until 2016 from the Survey of Consumer Finances (SCF) to measure financial choices of households. The SCF is a triennial repeated cross-section analysis sponsored by the Federal Reserve Board. It is carried out at the household level but collects individual demographic characteristics and income variables as well as detailed information on joint asset holdings of the household. To account for increased likelihood of survey non-response for asset-rich households, the SCF oversamples that population group. Since the focus of my study is not on the very rich, I drop the richest 10% of the sample (in terms of financial wealth). Dropping the upper decile of observations basically affects only couples and hence, does not change the empirical gender investment gap across single men and single women. In contrast, the resulting age profile for average financial wealth of couples is more comparable to measures from other datasets that do not oversample asset-rich households (such as the PSID). Moreover, to ensure the representativeness of the US population, I weigh each observation by the provided survey weight throughout the estimation.

For income variables, labor market outcomes and demographic characteristics I work with data from the Panel Study of Income Dynamics (PSID) spanning from 1989 until 2017 (Panel Study of Income Dynamics, 2021). The PSID is a longitudinal panel-survey of private households in the US running from 1968 until today.[2] Besides the core sample, the PSID oversamples low-income families (the 'SEO' sample) and immigrant families (the 'immigrant' sample). To make the sample comparable to that from the SCF, I drop all families belonging to those two sub-samples. Each wave, household members report biographic information, their individual labor force status and individual income levels. All financial variables are converted into 2007 dollars using the CPI-U.[3]

I define a single woman to be a family unit with a female head and no spouse present. Single men are defined accordingly. In total, the PSID sample consists of 57,986 individual-year observations that correspond to 701 single women, 593 single men and 4,050 individuals who live in couples. In contrast, the data drawn from the SCF includes information on 8,513 individuals in couples, on 950 single men and on 2,848 single women.

2.2.2 Life-Cycle Profiles of Portfolio Allocation

Throughout the analysis, I define financial assets in gross terms, that is financial wealth net of housing assets and debt (i.e. mortgages). Risky assets include direct stock holdings, corporate and foreign bonds, the fraction of mutual funds that include the former as well as the fraction of retirement accounts which is invested in stocks.[4]

Figure 2.1a displays the life-cycle profiles of equity shares for single men, single women and for couples during their working age. Equity shares combine the extensive margin (whether or not the households owns any risky assets) with the intensive margin (conditional on holding risky asset, what portfolio share is allocated to them?). To obtain a more complete picture, Figure 2.1b and Figure 2.1c separately plot the stock market participation rate (that is, only the extensive margin) and the conditional risky share (that is, only the intensive margin), respectively. I obtain all graphs by linearly regressing the respective dependent variable on a constant, age, the second polynomial of age, an interaction term of gender and age, a dummy that indicates more than 12 years of education, the number of household

[2] Because the Survey of Consumer Finances starts in 1989, I restrict my data sample taken from the PSID to the waves from 1989 until 2017. Data were collected annually until 1997 and afterwards every two years.

[3] CPI estimates taken from the US Bureau of Labor Statistics, available under this link [Accessed May 22, 2019].

[4] In Appendix B.1, I show that my results are robust to adopting a tighter definition of risky assets that does not include retirement accounts.

members, labor income and year dummies.[5] All differences are statistically different from zero, see Table 2.1 for the corresponding regression coefficients.

Figure 2.1 Life Cycle Pattern of Household Finances (Data)

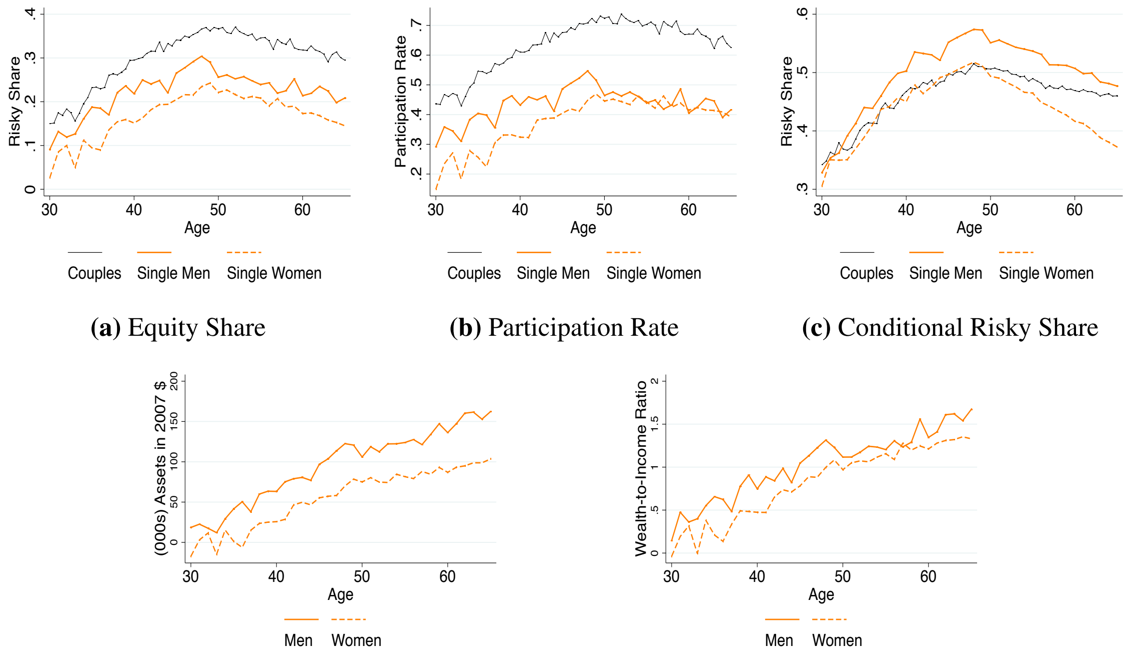

(a) Equity Share **(b)** Participation Rate **(c)** Conditional Risky Share

(d) Financial Assets (Singles) **(e)** Wealth-to-Income Ratio (Singles)

Notes: Figure 2.1 plots the life-cycle profiles of the equity share, stock market participation rates and conditional risky shares for singles and couples as well as absolute financial assets and the wealth-to-income ratio of singles. The sample consists of individuals born between 1945 and 1960 in the waves 1989 until 2016 of the Survey of Consumer Finances (SCF). Risky assets are defined as direct stock holdings, corporate and foreign bonds, the fraction of mutual funds that include the former and the fraction of retirement accounts which is invested in stocks.

Figure 2.1a shows that the equity share of single women is lower than that of single men during their entire working life. On average, the equity share of single women is 4.79%pts lower than that of men which – given an average equity share of single men of 23.43% – corresponds to 20.44% and roughly remains constant over the life-cycle. In contrast, the observed gender gap in stock market participation rates (Figure 2.1b) converges towards the entry to retirement whereas the gender gap in the conditional risky share diverges with age (Figure 2.1c).

Furthermore, the black solid line in Figure 2.1a shows that couples have on average a higher equity share than singles which is mainly driven by the extensive margin (see the black solid lines in Figures 2.1b and 2.1c, respectively). However, this finding is partly mechanical as couples are composed of two individuals for whom I compute the joint participation probability of participation. If I randomly draw a single men and a single women and compute the likelihood that at least one of them holds risky assets

5 To account for observations with zero or very little labor income, I transform labor income into its inverse hyperbolic sine before running the regressions.

(conditional on age), the participation rate of such a "generated couple" closely aligns with the ones of couples in the data.

Finally, Figure 2.1d confirms that single women accumulate less wealth than single men what is often referred to as the "gender wealth gap". Over the course of their working life, the gap in financial wealth is on average \$37,760, being largest when entering retirement (\$57,300). The gender wealth gap also prevails when normalizing by current labor income (as shown by the wealth-to-income ratio in Figure 2.1e).

2.2.3 Regression Coefficients over the Life-Cycle

The empirical gender differences in portfolio choices reported in Figure 2.1 can arise due to differences in circumstances or due to differences in preferences. The objective of this paper is to quantify the importance of the former. As a first exercise, I consider linear regressions that control for observable characteristics (Table 2.1). In particular, I run Tobit regressions (to account for non-participating households) of the equity share on a gender dummy, age polynomials and gender interacted with age (Column (1)). In Column (2), I additionally control for observable characteristics that the literature has shown to be important predictors for portfolio choices. Following Christelis et al. (2013), I control for the education of the individual, the overall number of household members and the inverse hyperbolic sine transformation of labor income. Finally, Column (3) furthermore includes the inverse hyperbolic sine transformation of the households' safe financial assets. However, since the amount of safe assets directly affects the equity share (that is defined the amount of risky assets over the sum of safe and risky assets), I treat Column (2) as the main specification throughout the rest of the paper. The corresponding marginal effects for the gender dummy along with their standard errors at various ages are reported in the last three rows of Table 2.1.[6]

The coefficient indicating whether the individual is a single women is negative (and statistically significant) across all three specifications. In contrast, the interaction term of gender and age is largest in the first column (least controls), slightly smaller in the second column and becomes statistically insignificant in the third specification. When considering the marginal effect of being a single woman on equity shares (see last three rows in Table 2.1, "ME"), I find a negative and significant gender effect across the entire life-cycle in the third specification that controls for the most observable characteristics.

6 Appendix B.2 lists the corresponding specifications for the participation rate (Table B.1) and for the conditional risky share (Table B.2).

In contrast, in the main specification (Column 2), the "negative" gender effect on equity shares disappears and eventually turns (slightly) positive as individuals age. Thus, the unexplained part of the gender investment gap (i.e. the part that is not accounted for when including controls) is strongest for young households. This remaining part of the gap can either arise because of unobserved heterogeneity across men and women (e.g. in preferences), because the mapping from observable characteristics to portfolio choices is non-linear or because I did not control for the correct observable characteristics.

Therefore, to further explore how differences in circumstances between single men and single women translate into heterogeneous portfolio choices over the life-cycle and to quantify their relative importance, Section 2.3 builds a structural model of gender and portfolio choice. Having a structural model helps to accommodate non-linearities and to account for factors that affect portfolio choices but cannot be easily controlled for in reduced form specifications, such as expectations and risk exposure.

2.2.4 On the (Non-)Presence of Housing

The focus of this paper is on liquid financial wealth which is why I abstract from housing wealth both in the empirical part as well as in the model (Section 2.3). However, in reality, housing constitutes a large share of households' portfolio and housing choices affect stock market behavior.[7] For the purpose of the current analysis, abstracting from housing is a problem if either housing choices directly map into portfolio behavior (and hence, the gender investment gap could be entirely explained by differences in housing) or if it differentially affects portfolio choices by gender, i.e. if housing is an important driver of the gender investment gap itself.

To explore whether either of these issues are present in the data, I conduct two exercises: First, if portfolio choices are a direct mapping of housing decisions, we would expect their life-cycle profiles to closely follow those in Figures 2.1a to 2.1c. Figure 2.2 displays the life-cycle profiles of singles for three different housing variables: the homeownership rate, gross housing wealth (henceforth: "HW") and the housing-wealth-to-income ratio (henceforth: "HI"). Single men hold on average slightly more housing wealth than single women. However, I do not find any meaningful gender differences in terms of homeownership rates or in the housing-wealth-to-income ratio despite significant gender differences in portfolio choices both along the extensive and intensive margin. Moreover, the life-cycle patterns for housing variables are different than those of portfolio choices: Neither housing graph displays a

7 One of the first papers to introduce housing in a model of portfolio choice were Cocco (2005) and Yao and Zhang (2005). Since then, there has been a large and ongoing literature on housing and portfolio choices, see for example Flavin and Yamashita (2011), Chetty et al. (2017) or Paz-Pardo (2020) to name a few.

hump-shaped pattern with a constant gender gap (as for the equity share), nor a converging gender gap (as for the stock market participation rate) nor a diverging gender gap (as for the conditional risky share).

Figure 2.2 Life Cycle Profiles of Housing Pattern (Singles)

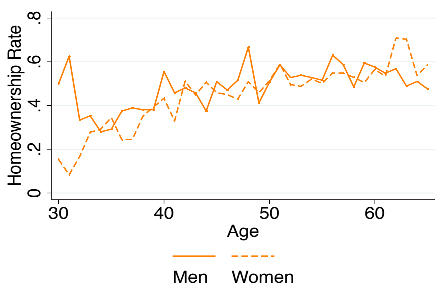

(a) Homeownership Rate

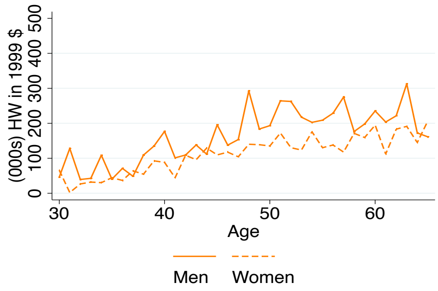

(b) (Gross) Housing Wealth

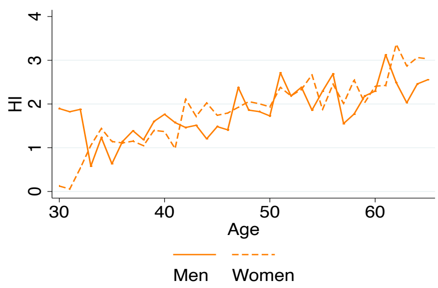

(c) Housing-Wealth-to-Income Ratio

Notes: Figure 2.2 plots life-cycle profiles of the homeownership rate, gross housing wealth and the housing-wealth-to-income ratio ("HI") for single men and single women. The sample consists of individuals born between 1945 and 1960 in the waves 1989 until 2016 of the Survey of Consumer Finances (SCF).

Second, to understand if housing differently affects the stock market behavior of single men and single women, I compare the predicted equity share and stock market participation rate of single homeowners to those of single non-homeowners, separately by gender. In particular, I first split the sample by housing tenure and run two separate regressions on stock market participation and equity shares, respectively, controlling for observable characteristics. Figure 2.3 plots the predicted outcome variable from these regressions for an individual with more than 12 years of schooling, the median income of singles, the median number of children for singles (zero) and who is at the respective age in 2001 (which is approx. the midpoint of my sample).

I find that generally, homeownership matters for predicted portfolio choices, in line with previous literature. However, albeit different in levels, the gender differences in predicted participation rates and equity shares are very similar for homeowners and non-homeowners (i.e. the differences between black and the orange line), especially during young age. The predicted gender gap in equity shares is slightly larger among home-owners towards the end of the life-cycle whereas the gap in participation rates and the gap in equity shares for young household does not significantly differ by homeownership-status, increasing my confidence that excluding housing from the analysis does not change the results regarding the sources of the gender investment gap while at the same time keeping the model tractable and easing computational complexity.

Figure 2.3 Life Cycle Profiles of Housing Pattern (Singles)

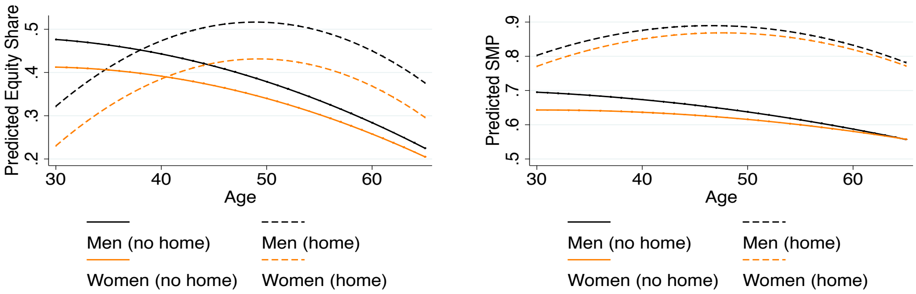

(a) Equity Share **(b)** Participation Rage

Notes: Figure 2.3 plots the predicted life-cycle profiles of the equity share and the stock market participation rates of a single individual in 2001 who has a high school degree, no children and the medium level of income and safe assets, separately by gender and housing tenure.

Table 2.1 Regression Coefficients & Marginal Effects – Equity Shares of Singles

	(1) Equity Share	(2) Equity Share	(3) Equity Share
single woman	-0.367***	-0.150***	-0.0333**
	(0.0153)	(0.0166)	(0.0165)
single woman*age	0.00544***	0.00159***	-0.000253
	(0.000301)	(0.000307)	(0.000267)
age	0.0768**	0.0118	0.00315
	(0.0330)	(0.0440)	(0.0582)
$age^2 * 100$	-0.0538	0.0134	0.0246
	(0.0710)	(0.0940)	(0.125)
$age^3 * 10000$	-0.0302	-0.0312	-0.0411
	(0.0495)	(0.0646)	(0.0858)
High education		0.368***	0.182***
		(0.00560)	(0.00589)
No. of HH members		-0.0511***	-0.0354***
		(0.00293)	(0.00370)
Income		0.0422***	0.0270***
		(0.00118)	(0.00103)
Safe assets			0.0853***
			(0.00147)
Constant	-2.088***	-1.331**	-1.607*
	(0.495)	(0.646)	(0.856)
Observations	4,737	4,735	4,735
Year FE	No	Yes	Yes
ME for women at age 30	-0.0405***	-0.0539***	-0.0485***
	(0.00497)	(0.00380)	(0.00238)
ME for women at mean age (50)	0.1817***	0.0113	-0.0588***
	(0.0160)	(0.0150)	(0.0109)
ME for women at age 65	0.340***	0.0578**	-0.0661***
	(0.0246)	(0.0238)	(0.0186)

Notes: Estimations are based on Tobit regressions on the sample of individuals that live in households with no spouse present. Source: SCF waves 1989 until 2016; individuals born between 1945 and 1960. Equity Share = Unconditional risky share. *single woman* is a dummy indicating that the household head is a women. *high education* is a dummy equal to one if the household head has more than 12 years of education. *safe assets* refers to safe liquid assets. "ME" indicates the marginal effect of being a women at the respective age. Robust standard errors in parentheses, *** $p<0.01$, ** $p<0.05$, * $p<0.1$

2.3 A Life-Cycle Model of Portfolio Choice

In the following, I develop a life-cycle model of portfolio choice along the lines of Cooper and Zhu (2016). I extend their set-up by introducing three types of households in the model: single men, single women and couples and allow for marital transitions across the life-cycle. In contrast, I abstract from adjustment costs of stock holdings.

2.3.1 Environment

In the model, agents can be women or men (denote gender by $i = \{f, m\}$) and live either as singles ($\mathcal{S}$) or as a married couple ($\mathcal{M}$). Thus, there exist in total four types of agents: single women ($\mathcal{S}, f$), single men ($\mathcal{S}, m$), married women ($\mathcal{M}, f$) and married men ($\mathcal{M}, m$). For all, their life can be split in two stages: working age and retirement. Time is discrete. A model period is one year long. Agents start their life at age 30, retire at age 65 and die deterministically at age 85, i.e. $j \in \{30, 31, ..., 65, ..., 85\}$. They face uncertain survival during retirement that depends on their age j. At age 30, agents are ex-ante heterogeneous in terms of their education θ which can take two values ($\theta = \{l, h\}$) and refers to college and non-college educated individuals in the data. I treat θ as exogenous and assume that agents enter the model after completing education.

During working age, when being single, individuals decide how much to consume (c_i) and how much to save in a safe asset (a_i^s) as well as how much to save in a risky asset (a_i^r). Couples decide jointly on the level of consumption ($c_{\mathcal{M}}$) as well as on how much to save in both types of assets ($a_{\mathcal{M}}^s, a_{\mathcal{M}}^r$). That is, consumption is treated as a public good and becomes private only upon divorce. Moreover, singles face an exogenous marriage probability each period that depends on their gender, age and education level. Likewise, couples face an exogenous divorce probability that again varies by age and both spouses' education.

During retirement, agents do no supply labor but receive a fixed pension which is a fraction of their last realized labor income. Additionally, they face age- and gender dependent medical expenditures and are subject to longevity risk. Upon dying, agents value leaving bequests. As during working age, they can live both as single or couple, however their marital status is fixed (i.e. there is no marriage or divorce). If one spouse living in a couple dies, the surviving spouse continues his or her life as a single with a fraction of the couples' assets to account for increased medical expenditures in the year prior to death as well as for bequests to non-spousal heirs (Jones et al., 2020). As before, agents have a portfolio choice between a safe asset and a risky asset.

2.3.2 Preferences

While I do not allow for gender heterogeneity in preferences, I introduce preference heterogeneity by marital status. In particular, I allow the discount factor β, the coefficient of relative risk aversion γ and the stock market participation costs S^F to vary between singles and couples. Empirically, couples have a higher savings rate than singles albeit their overall income variance being lower, contradicting model predictions. The higher savings rate can arise from various sources, such as saving for children's college, higher homeownership rates among couples or differences in preferences. As the focus of this paper is to explain gender heterogeneity in investment pattern and not differences across singles and couples, I choose to introduce preferences heterogeneity by marital status that allows me to accommodate the data while keeping the model tractable and focusing on the core research question.

Singles. Single individuals can either be a man or a woman with their gender being denoted by $i = \{f, m\}$. They have time-separable CRRA preferences over a consumption good c_i. The period flow of utility is given by:

$$u(c_i) = \frac{\eta_{ij}\left(\frac{c_i}{\eta_{ij}}\right)^{1-\gamma_s}}{1-\gamma_s}$$

where γ_s is the coefficient of relative risk aversion that is fix across gender and η is an equivalence scale that adjusts for household size and which is allowed to vary by age j and gender i.

Couples. Each couple is composed of exactly one woman and one man. As for singles, couples have time-separable CRRA preferences over the consumption good $c_{\mathcal{M}}$ which is public within the household. Their period flow of utility can therefore be expressed as:

$$u(c_{\mathcal{M}}) = \frac{\eta_{cj}\left(\frac{c_{\mathcal{M}}}{\eta_{cj}}\right)^{1-\gamma_c}}{1-\gamma_c}$$

Again, γ_c is the coefficient of relative risk aversion (which can be different to that of singles, γ_s) and η is an age-dependent equivalence term adjusting for household size.

Bequest Motive. In the event of death, individuals derive utility from leaving bequests according to:

$$\phi(a') = L\frac{(\omega + a')^{1-\gamma}}{1-\gamma}$$

where a' denotes the be-quested assets, ω captures the luxuriousness of the bequest motive and L governs the bequest intensity. Bequest preferences are homogeneous across all types of households. Couples value leaving bequests only if they both die within the same period. Whenever only one spouse dies, the surviving spouse continues life as a single with a fraction of the couples' assets (and hence, values leaving bequest in the case of his or her own death).

2.3.3 Dynamics

Asset Returns. Agent accumulate savings for retirement and to smooth consumption. To do so, they have access to two types of assets: One safe and one risky asset, denoted by a_s and a_r, respectively. The safe asset pays a time-invariant return r_s. In contrast, the return of the risky asset is drawn from the distribution $r_r \sim N(\mu_r, \sigma_r^2)$ that is assumed to be i.i.d and for which it holds that $\mu_r > r_s$. Following Fagereng et al. (2017), I allow for the possibility of stock market crashes and augment the return of the risky asset by a "disaster" state. That is, with probability $(1 - p_{tail})$ the return is drawn from the above normal distribution and with probability p_{tail} a tail event $r_{tail} < \underline{r}_r$ materializes. Short-selling and borrowing are not allowed.

Income Profiles. Following Borella et al. (2019), I assume that income can be split into a deterministic and into a stochastic component. More precisely, income y_{ij} at age j for gender i can be expressed as:

$$y_{ij} = \bar{y}_i \theta_i \xi_{ij} \tilde{y}_{ij}$$

where $\bar{y}_i$ denotes a constant, θ_i is the (exogenous) education premium and ξ_{ij} stands for an age-specific component. Finally, $\tilde{y}_{ij}$ represents the stochastic component of income consisting of a transitory and a persistent shock:

$$\tilde{y}_{ij} = z_{ij} + \epsilon_{\tilde{y}ij}$$

$$z_{i,j+1} = \rho_{zi} z_{ij} + \nu_{zij}$$

where $\epsilon_{\tilde{y}ij}$ and ν_{zij} are independent zero mean random shocks with variances $\sigma^2_{\tilde{y}i}$ and σ^2_{zi} respectively. The parameter $\rho_{zi} \in (0, 1]$ captures the persistence of shock ν_{zi}.

All parameters of the income process are allowed to vary by gender and by marital status to account for the fact that marriage typically results in lower income for women whereas it increases earnings for men.[8] Within couples, the transitory shocks ν_{zfj} and ν_{zmj} are allowed to be correlated as spouses live in the same area and are likely to work in similar industries (family business, they meet at work) and are thus subject to correlated labor market shocks. In contrast, following Cocco et al. (2005), labor income shocks are uncorrelated to realizations of the stock return.

Out-of-Pocket Medical Expenditures. When being retired, agents are subject to medical expenditures m_j that are a deterministic function of age and gender. However, because individuals face survival risk and because medical expenditures are strictly increasing in age, deterministic medical expenditures impose a source of risk in the sense that agents are uncertain whether or not they live until a certain age and have to pay the corresponding medical bills. This modeling choice is motivated by De Nardi et al. (2010) who show that the main sources of risk during retirement are not fluctuations of medical expenditures around its mean but rather their age-dependent level combined with longevity risk.

2.3.4 Stock Market Participation Cost

In order to avoid the model to predict excess stock market participation rates, agents have to pay a fixed cost S^F each period if they choose to invest part of their savings in the risky asset. This cost is allowed to differ between couples and singles, however, it is equal for single men and single women. As in Vissing-Jorgensen (2002), I model participation costs as a flow variable, that is they have to be paid each period irrespectively of the history of stock holdings. The main advantage to model participation costs as a flow variable rather than an entry cost (see e.g. Alan (2006) or Cooper and Zhu (2016)) is that flow costs do not require introducing stock holdings as an additional state variable and therefore reduce the computational complexity of the model which is – considering the different household structures – already quite substantial.

[8] The empirical observed "marriage penalty" for women's earnings can arise because of self-selection into marriage (i.e. low income women are more likely to get married) or because marriage itself affects female income, e.g. through childbirth or household specialization. Analyzing the relative importance of these factors is beyond the scope of this paper and therefore, I impose women's earnings drop upon marriage exogenously.

2.3.5 Marriage and Divorce

Single individuals get married with an exogenous probability that depends on their gender, their age, their education and on their current productivity realization. Denote this marriage probability by $\mu_i(j, \theta, \tilde{y})$. Conditional on meeting a partner, the probability of meeting a partner with education θ_p and shock realization $\tilde{y}_p$ is:

$$\Pi(.) = \Pi(\theta_p, \tilde{y}_p | \theta_i, \tilde{y}_i)$$

Both partners always have the same age. Individuals are always matched to a partner with the mean empirical amount of assets (conditional on age, gender and education). This specification generates assortative mating along asset holdings as we observe it in the data while at the same time allowing for the possibility that within couples, the income of the husband is usually higher than that of the wife. Couples face an exogenous divorce probability each period that depends on age, the education of each spouse as well as both productivity realizations, that is the likelihood of divorce can be expressed as $\lambda(j, \theta_f, \theta_m, \tilde{y}_f, \tilde{y}_m)$. Upon divorce, assets are split equally between spouses and 25% of assets are destroyed to account for legal fees of divorce and general costs of splitting assets between spouses.

2.3.6 Timing

Timing within one period is as follows. In the beginning of period t all shocks materialize. That is, agents learn their current productivity state(s), their stock market return as well as their marital status. Thus, agents start period t with a given amount of savings that depends on their decisions in period $t-1$, their marital status and the realization of the asset return state. After observing all shock realizations, agents decide on how much to consume and how much to save in both the risky and the safe asset. When investing part of their endowment in the risky asset (i.e. if $a_{r_{t+1}} > 0$), they to pay S^F in the *current* period, that is in period t.

2.3.7 Recursive Formulation

I express the problem recursively by defining six value functions: the value function for singles, the value function for couples and the value function for an individual living in a couple, all during working age as well as during retirement. The latter is the relevant object when computing the present value of marriage for a single whereas the value function for couples determines the optimal allocation of resources within a couple across time (Borella et al., 2019). Moreover, because the stock market participation cost has to be paid per-period and given the i.i.d nature of the return process for the risky asset, I can combine safe and risky assets into one "asset cash-in-hand" state variable: $a = (1+r_r)a_r + (1+r_s)a_s$.

Singles – Working Age. The state variables of a single agent are her gender i, age j, education θ, asset cash-in-hand a and her current income realization $\tilde{y}$. Each period, she has a consumption-savings choice and additionally decides on how to split her savings between the safe and the risky asset. The corresponding value function reads as:

$$V^S(i,j,\theta,a,\tilde{y}) = \max_{a'_s \geq 0, a'_r \geq 0, c \geq 0} \frac{\eta_{ij}\left(\frac{c}{\eta_{ij}}\right)^{1-\gamma_s}}{1-\gamma_s} + (1-\mu(j,\theta,\tilde{y}))\beta_s s \mathbb{E}V^S(i,j+1,\theta,a',\tilde{y}') +$$
$$\mu(j,\theta,\tilde{y})\beta_s \mathbb{E}\hat{V}^C(i,j+1,\theta,\theta_p,a'+a'_p,\tilde{y}',\tilde{y}'_p)$$

subject to:

$$a'_r + a'_s + c = y(j,\theta,\tilde{y}) + (1+r_s)a_s + (1+r_r)a_r - \mathbb{1}_{a'_r>0}S^F_s$$

$$a = (1+r_r)a_r + (1+r_s)a_s$$

and:

$$\tilde{y} = z + \epsilon_{\tilde{y}} \quad \text{with} \quad z' = \rho_z z + \nu_z \quad \text{and} \quad \epsilon_{\tilde{y}} \sim N(0,\sigma^2_{\tilde{y}}), \nu_z \sim N(0,\sigma^2_z)$$

$$r_r \sim N(\mu_r,\sigma^2_r) \quad \text{with} \quad \mu_r > r_s \quad \text{and} \quad \tilde{y} \perp r_r, \quad \mathbb{E}(\tilde{y}',r'_r,\Pi|j,\theta,z)$$

where η_j denotes an equivalence parameter that controls for changing family size over the life-cycle. $\hat{V}^C$ expresses the value of individual i of getting married to partner p. Single individuals take the expected value over future productivity realizations and asset return when staying single whereas they form expectations over future productivity realization, asset returns and their specific partner in case of getting married.

Singles – Retirement. During retirement, agents do not supply labor and receive a fixed pension income that depends on their last labor income realization. There is no marriage or divorce. Retired individuals face survival risk. Moreover, they are subject to deterministic age-dependent medical expenditures m_{ij}, leaving as state variable gender i, age j, education level θ, asset cash-in-hand a as well as the last income realization before retirement ($\hat{y}$). Each period, retired singles face a consumption-saving as well as a portfolio choice and they value leaving bequests in the case of death.

$$V_R^S(i,j,\theta,a,\hat{y}) = \max_{a'_s \geq 0, a'_r \geq 0, c \geq 0} \frac{\eta_{ij}\left(\frac{c}{\eta_{ij}}\right)^{1-\gamma_s}}{1-\gamma_s} + \beta_s \psi_{ij} \mathbb{E} V_R^S(i,j+1,\theta,a') + \beta_s(1-\psi_{ij}) L \frac{(\omega + a')^{1-\gamma}}{1-\gamma}$$

subject to:

$$a'_r + a'_s + c = pen(\hat{y}) + (1+r_s)a_s + (1+r_r)a_r - m_{ij} - \mathbb{1}_{a'_r > 0} S_s^F$$

$$a = (1+r_r)a_r + (1+r_s)a_s$$

$$r_r \sim N(\mu_r, \sigma_r^2) \quad \text{with} \quad \mu_r > r_s, \quad \text{and} \quad \mathbb{E}(r'_r)$$

where ψ_{ij} denotes the age-dependent survival probability that differs between men and women. Retired singles take the expected value over their next-period asset return as well as their likelihood of survival.

Couples – Working Age. The value function for couples during working age is needed to compute optimal allocation for a couple that consists of a women f and a man m. While I allow for both marriage and divorce during working age, individuals cannot switch partners between two consecutive periods. The state variables of a couple can be summarized by their age j (which is assumed to be the same), education of both spouses θ_f, θ_m, their joint asset holdings a as well as both current productivity realizations $\tilde{y}_f, \tilde{y}_m$. The corresponding value function reads as:

$$\begin{aligned} V^C(j,\theta_f,\theta_m,a,\tilde{y}_f,\tilde{y}_m) = \max_{a'_s \geq 0, a'_r \geq 0, c \geq 0} & \frac{\eta_{\mathcal{M}j}\left(\frac{c}{\eta_{\mathcal{M}j}}\right)^{1-\gamma_c}}{1-\gamma_c} + \\ & (1-\lambda(j,\theta_f,\theta_m,\tilde{y}_f,\tilde{y}_m))\beta_c \mathbb{E} V^C(j+1,\theta_f,\theta_m,a',\tilde{y}'_f,\tilde{y}'_m) + \\ & \lambda(j,\theta_f,\theta_m,\tilde{y}_f,\tilde{y}_m)\beta_c \sum_{i=f,m} \mathbb{E} V^S(i,j+1,\theta_i,0.75\frac{a'}{2},\tilde{y}'_i) \end{aligned}$$

subject to:

$$a'_r + a'_s + c = \sum_{i=f,m} y(j,\theta_i,\tilde{y}_f,\tilde{y}_i) + (1+r_s)a_s + (1+r_r)a_r - \mathbb{1}_{a'_r > 0} S_c^F$$

$$a = (1+r_r)a_r + (1+r_s)a_s$$

and:

$$\tilde{y}_i = z_i + \epsilon_{\tilde{y}_i} \quad \text{with} \quad z_i' = \rho_{zi} z_i + \nu_{zi} \quad \text{and} \quad \epsilon_{\tilde{y}_i} \sim N(0, \sigma^2_{\tilde{y}_i}), \nu_z \sim N(0, \sigma^2_{zi}) \quad \text{for} \quad i = \{f, m\}$$

$$\begin{pmatrix} \nu_{zf} \\ \nu_{zm} \end{pmatrix} \sim \begin{pmatrix} \sigma^2_{zf} & \rho_{\sigma_{zf}, \sigma_{zm}} \\ \rho_{\sigma_{zf}, \sigma_{zm}} & \sigma^2_{zm} \end{pmatrix}$$

$$r_r \sim N(\mu_r, \sigma_r^2) \quad \text{with} \quad \mu_r > r_s, \quad \tilde{y} \perp r_r \quad \text{and} \quad \mathbb{E}(\tilde{y}_f', \tilde{y}_m', r_r' | j, \theta_f, \theta_m, \tilde{y}_f, \tilde{y}_m)$$

Couples take the expected value of both partners' future productivity realizations and joint asset returns when staying married as well as the respective individual's productivity realization and asset return when getting divorced. Moreover, the transitory parts of the income processes (ν_{zf} and ν_{zm}) are allowed to be correlated within couples.

Couples – Retirement. Retired couples receive a flat pension income that depends on the man's last income realization before retirement ($\hat{y}_m$). They do not work and cannot get divorced. However, they individually face the risk of dying. If one spouse dies, the surviving one continues his or her life as single with a fraction δ_i of the couple's assets. If both spouses die within the same period, they jointly value leaving bequests. Their value function reads as:

$$\begin{aligned} V_R^C(j, \theta_m, a, \hat{y}_m) = & \max_{a_s' \geq 0, a_r' \geq 0, c \geq 0} \frac{\eta_{\mathcal{M}j} \left(\frac{c}{\eta_{\mathcal{M}j}} \right)^{1-\gamma_c}}{1-\gamma_c} + \beta_c \psi_{jf} \psi_{jm} \mathbb{E} V_R^C(j+1, \theta_m, a', \hat{y}_m) + \\ & \beta_c \sum_{i=f,m} \psi_{ij} (1 - \psi_{-ij}) \mathbb{E} V_R^S(i, j+1, \theta_m, \delta_i a', \hat{y}_m) + \\ & \beta_c (1 - \psi_{jf})(1 - \psi_{jm}) L \frac{(\omega + a')^{1-\gamma}}{1-\gamma} \end{aligned}$$

subject to:

$$a_r' + a_s' + c = pen_c(\hat{y}_m) + (1 + r_s) a_s + (1 + r_r) a_r - \sum_{i=f,m} med_{ij} - \mathbb{1}_{a_r' > 0} S_c^F$$

$$r_r \sim N(\mu_r, \sigma_r^2) \quad \text{with} \quad \mu_r > r_s, \quad \text{and} \quad \mathbb{E}(r_r')$$

Thus, retired couples take the expected value over their joint asset return as well as the individual's survival probabilities.

Value to an individual of becoming a couple. The value of an individual in a couple is the relevant object when computing the value of single i for getting married to partner p, i.e. the present discounted value of the individual's utility in the event of marriage (Borella et al., 2019). In this context, variables denoted with a $\hat{hat}$ indicate optimal allocations computed with the value function for couples, given the respective state variables. The value of an individual in a retired couple $\hat{V}_R^C$ is defined accordingly.

$$\hat{V}^C(i,j,\theta_i,\theta_p,a,\tilde{y}_i,\tilde{y}_p) = \frac{\eta_j \left(\frac{\hat{c}}{\eta_j}\right)^{1-\gamma_c}}{1-\gamma_c} + (1-\lambda(j,\theta_i,\theta_p,\tilde{y}_i,\tilde{y}_p))\beta_c \mathbb{E}\hat{V}^C(i,j+1,\theta_i,\theta_p,a',\tilde{y}_i',\tilde{y}_p') + \lambda(j,\theta_i,\theta_p,\tilde{y}_i,\tilde{y}_p)\beta_c \mathbb{E}V^S(i,j+1,\theta_i,\frac{a'}{2},\tilde{y}_i')$$

2.4 Estimation

As in Gourinchas and Parker (2002) or Cagetti (2003), I estimate the model using a two-step strategy. That is, I first estimate all parameters that can be cleanly identified outside of the model and pre-set some parameters to values from the literature. In a second step, I estimate the remaining structural parameters using the Simulated Method of Moments (SMM), taking the parameters from the first stage as given. First stage parameters include initial distributions, parameters related to medical expenditures, the labor income process, survival probabilities and asset returns. I borrow the parameters for the bequest motive (ω, L) from Cooper and Zhu (2016) who estimate bequest parameters in a portfolio choice context with CRRA preferences. Consequently, second stage parameters include the discount factor β, the coefficient of relative risk aversion γ as well as the stock market participation cost S_s^F, all separately for singles and couples. I collect the second stage parameters in the vector $\Theta = \{\beta_s, \beta_c, \gamma_s, \gamma_c, S_s^F, S_c^F\}$.

2.4.1 First Stage Estimation

Income Profiles. For most individuals, their ability to accumulate wealth and the decision on how to invest that wealth is strongly affected by their life-cycle profile of income. In particular, women typically have lower income than men, leading to heterogeneous financial choices. This gender discrepancy is especially pronounced within couples. Figure 2.4 shows life-cycle profiles of average income by gender and by marital status. Income is expressed as annual income out of labor earnings (including labor income from farms and businesses), social security benefits and transfers.When estimating those profiles, I restrict the sample to individuals who did not change their education after age 30 because education is exogenous in the model. Moreover, for singles, I include labor earnings, social security benefits and transfers from all members of the households to ensure that my measure of income adequately accounts for disposable income of single households in the data. For couples, I assign each spouse their own labor income, social security benefits and transfers and add half of that from other household members.[9] Lastly, I drop observations who, according to the described measure, report zero annual income (in the case of couples, if they report zero overall income).

To estimate the income profiles, I follow Borella et al. (2019) and first split the sample by marital status and then separately regress the log of income for individual i at age j,

$$ln(income_{ij}) = \alpha + \beta_1 age_{ij} + \beta_2 age_{ij}^2 + \beta_3 woman_i * age_{ij} + \beta_4 woman_i * age_{ij}^2 + \delta_i + u_{ij}$$

on a fixed effect δ_i, age, age^2 as well as their interaction term with a dummy that indicates if the individual is a woman. To obtain shifters for both gender and education level, I regress the sum of the fixed effect and the residual on fully interacted dummies of gender and education level:

$$\delta_i + u_{ij} \equiv w_i j = \gamma_0 + \gamma_1 woman_i + \gamma_2 educ_i + \gamma_2 woman_i * educ_i + \epsilon_{ij}$$

where $educ_i$ is defined as a dummy taking the value one if the respective individuals has more than 12 years of schooling.

Figure 2.4 Life Cycle Profiles of Income (Deterministic Component)

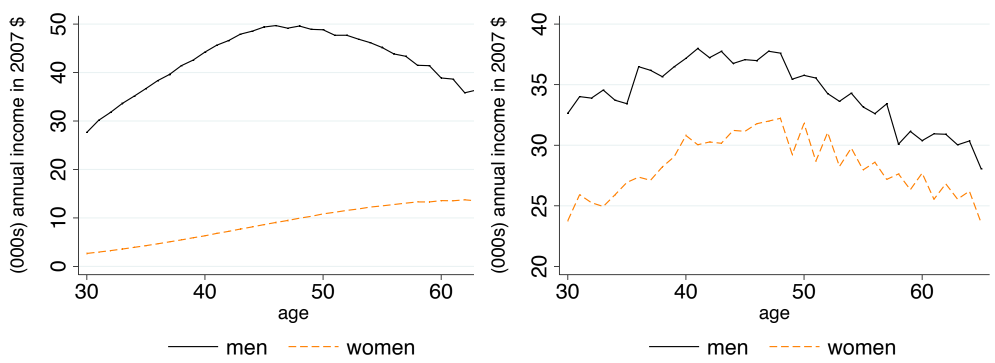

(a) Couples **(b)** Singles

The coefficients from these income equations (reported in Table B.3 in Appendix B.3.2) inform me about the deterministic component of the income process in the model which can be split into a constant,

9 For some years, the PSID does not separately report transfer income or social security benefits of spouse and household head. In these cases, I allocate half of the overall reported measure to the wife and the other half to the husband.

an exogenous education dummy and an age-specific part. I estimate the parameters governing the stochastic component of the income process using the minimum distance estimator as in Guvenen (2009).[10] Table 2.2 summarizes the results. My point estimates imply a slightly less persistent income process for single women than for single men and than for married individuals. Moreover, the variance of the persistent shock σ_z^2 is higher for single women than for the rest of the population. Thus, single women face overall a more risky income process than single men, however, the difference is relatively small in magnitude. Notably, the income process of married women exhibits a much higher variance of the transitory shock $\sigma_{\tilde{y}}^2$ than that of singles and that of married men.

Table 2.2 Estimation Results – Stochastic Income Process

Parameter	Men	Women	Men	Women
	Singles		Couples	
ρ_z	0.937	0.9138	0.9307	0.9369
	(0.0131)	(0.0173)	(0.0065)	(0.0048)
σ_z^2	0.087	0.1007	0.0815	0.0882
	(0.021)	(0.0236)	(0.0081)	(0.0097)
$\sigma_{\tilde{y}}^2$	0.1269	0.1175	0.0928	0.2949
	(0.0625)	(0.0417)	(0.0206)	(0.0319)

Notes: Standard Errors in parentheses obtained with bootstrapping (2000 replications).

Marital Transitions. Besides income fluctuations, one main source of financial risk during working age is marital status and possible changes thereof. Figure A.13a plots marriage probabilities by age, gender and education whereas Figure A.13b displays divorce probabilities by education and age. Both graphs are estimated using PSID data. Marital transitions are defined as the likelihood of getting married (respectively divorced) within the next period conditional on not being married (respectively being married) in the current period. More specifically, I estimate the following logit function, separately for couples and singles:

$$\xi_{t+1} = \frac{exp(X_t\beta^s)}{1+exp(X_t\beta^s)}$$

where ξ_{t+1} denotes the probability of being married (respectively divorced) next period. As explanatory variables, I include the age, age-squared, a dummy indicating whether the individual has some college education as well as a dummy for waves after 1997 to account for switch from annual to biannual frequency in the PSID.[11] Table B.4 in Appendix B.3.3 reports the corresponding regression coefficients.

[10] Details on the estimation strategy for the stochastic part of the income process can be found in Appendix B.3.1.

[11] For couples, all demographic variables refer to the household head.

Figure 2.5 Marital Transition Probabilities

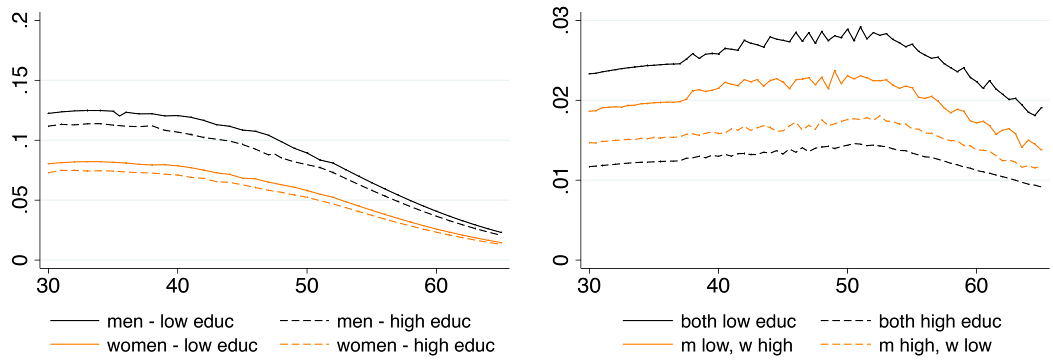

(a) Marriage Probability

(b) Divorce Probability

I find that marriage probabilities are higher than divorce probabilities, especially for young individuals. And any given age, single women are less likely than single men to get married within the next year. Moreover, the likelihood of divorce displays a hump-shaped pattern whereas the hazard of marriage declines over the life-cycle. The probability of divorce is decreasing in the education of both spouses (Figure A.13b). However, couples in which the husband has low education and the wife has a high education are more likely to get divorced than couples whose education is allocated in the opposite way. In contrast, the likelihood of getting married does not significantly differ by education (Figure A.13a).

Moreover, I estimate the marriage market (Π) non-parametrically directly from PSID data. Given that marriage occurs exogenously in the model, it may happen that individuals are matched to a partner although they had endogenously chosen to remain single. However, in almost 85% of cases, individuals prefer marriage over singlehood because marriage offers income pooling as well as economies of scale for consumption.

Out-of-Pocket Medical Expenditures. Upon entering retirement, households receive a flat pension income, eliminating their exposure to uninsurable income fluctuations. In contrast, they are subject to out-of-pocket medical expenditures that sharply increase towards the end of the life-cycle. In the model, those expenditures are assumed to be deterministic. However, given that individuals face uncertain survival, medical expenditures impose an uncertainty on households in the sense that it is unclear whether or not the individual survives up to that age. I borrow the parameters describing medical expenditures by age and gender from Borella et al. (2019). The authors estimate deterministic out-of-pocket medical expenditures profiles with data from the HRS separately for men and women who were born in the 1950s. They estimate higher medical expenditures for men at the start of retirement but a steeper gradient for women, especially after age 76. Moreover, to account for the possibility of informal care arrangement among spouses, I assume that medical expenses for married individuals are 80% than that of singles.

Survival Probabilities. I take gender specific death probabilities from the Life Tables of the US Social Security Administration.[12] The death probability at age j is defined as the probability to die within the next year conditional on having survived up to age j. I compute the inverse of those probabilities and work with average values between the years 1990, 2000 and 2010, corresponding to the sample period of my study. In the case of couples, both spouses face individual survival risk and thus, they may die in separate years. If the husband dies, the surviving wife keeps 60% of the household's assets, whereas a surviving husband keeps 70% of the household's asset to account for sharply increasing medical expenses in the year prior to death as well as for bequests to non-spousal heirs (Jones et al., 2020).

Asset Returns. I set the annual return rate of the risk-free asset to 3%. The risky asset has a normally distributed return plus a tail risk. That is, I first assume a risk premium of 3%, and a variance of $Var(\tilde{R}(s)) = \sigma_r^2 = (0.1758)^2$. The latter reflect the variance of the annual total return index of the S&P 500 from 1989 until 2016. Next, I augment the return of the risky asset by a tail event as in Fagereng et al. (2017). In particular, the return of the risky asset is drawn with probability $1 - p_{tail}$ from the normal distribution and with probability p_{tail} a disaster state materializes. p_{tail} is set to 2% and results in a loss of 50% of all risky assets. A disaster state accounts for severe stock market crashes that we observe in the data but that would not be accounted for when approximating the risky asset return by a normal distribution. Moreover, introducing negative skewness lowers the propensity of agents in the model to invest in risky assets and thus helps to resolve the issue that standard portfolio choice models typically predict excessive equity shares when compared to the data.[13] In the model, the asset return realization is an aggregate shock. When simulating the model for a large set of individuals born between 1945 until 1960 over their life-cycle, I simulate the return of the risky asset to mimic the observed stock market performance in the US when the cohort was in that respective age. In particular, I assume that 20% of the simulated sample were born between 1945 and 1948, 20% between 1949 and 1952 and so on.

[12] All tables available under this link [Accessed May 14, 2019].

[13] An alternative approach to generate lower equity shares is to introduce adjustment costs for the risky asset (see e.g. Cooper and Zhu (2016)). However, adjustment costs require to introduce the equity share as a state variable. To keep the model tractable, I therefore abstract from these adjustment costs.

Pension Payments. Pension payments are flat and assumed to be 60% of the income during the last year of work. That is, pensions differ by education and productivity state at age 65. Couples receive a common pension which is 1.7 times higher than that of single men.

Equivalence Scales. In the model, the equivalence scales η are allowed to differ by age and household structure (i.e. single men, single women, couple). To compute them, I first estimate the average household size by age and household structure from the PSID and then apply the OECD equivalence scale: I assign a weight of 1 to the first adult household member, a weight of 0.7 to all other adult member and a weight of 0.5 to each child.

Initial Conditions. The initial distribution over asset holdings in the model is chosen such that it mimics the distribution of wealth across individuals born between 1945 and 1960 at age 30 in the SCF. Similarly, I set the fraction of high and low educated individuals by gender to be the average share of individuals with more respectively less than 12 years of schooling in the PSID of that cohort. Finally, the initial distribution of couples and singles is set equal to PSID data for individuals at age 30 born between 1945 and 1960.

2.4.2 Second Stage Estimation

Taking the parameters from the first stage as given, I estimate the remaining structural parameters $\Theta = \{\beta_c, \beta_s, \gamma_c, \gamma_s, S_c^F, S_s^F\}$ using the Simulated Method of Moments. The exercise is to find $\hat{\Theta}$ that solves the following optimization problem:

$$\mathcal{L} = \min_{\Theta}(M^s(\Theta) - M^d)W(M^s(\Theta) - M^d)' \tag{2.1}$$

where W represents a weighing matrix, M^d moments derived from the data and $M^s(\Theta)$ their theoretical counterparts derived from model simulations.

Parameter Identification & Choice of Moments. The key challenge is to separately identify the coefficient of relative risk aversion γ, the discount factor β and the stock market participation cost S^F because all moments directly affect savings behavior and portfolio choices of households. Hence, different parameter values are not entirely orthogonal to one another which makes their separate identification difficult. In this section, I provide (informal) intuition why my moments of choice are informative about the parameters in question. Once households cross the threshold of participation, the participation cost S^F becomes irrelevant for their decision on how much to invest in the risky asset.

Taking this discrepancy into account, I identify γ by exploiting heterogeneity in the portfolio share across *participating* households, that is, in the conditional risky share. Moreover, I use heterogeneity in wealth levels to identify β and S^F to match the life-cycle profiles of participation rates. In particular, I target the life-cycle profiles for the conditional risky share, for the participation rate and for absolute wealth levels of couples to identify parameters referring to couples in the model $(\beta_c, \gamma_c, S_c^F)$. For parameters referring to singles $(\beta_s, \gamma_s, S_s^F)$, I target the corresponding life-cycle profiles of single men. Consequently, life-cycle profiles of single women serve as untargeted moments to validate the model.

The Weighting Matrix W. I first estimate the 2nd stage parameters by using the identity matrix, i.e. with $W = \mathcal{I}$. Consequently, every moment receives equal weight in the estimation procedure. In a second run, I use the inverse of the variances of my moment conditions as a (diagonal) weighting matrix in order to assign a lower weight to less precisely estimated data moments ($W = \frac{1}{\mathcal{V}}$). This approach follows Cooper and Zhu (2016) and is in contrast to most papers that use the standard variance-covariance matrix (e.g. Cagetti (2003) or Alan (2006)). However, in the current set-up, different moments are based on different sample sizes: While the participation rates and wealth levels include all observations, the conditional risky share only includes stock market participants. Hence, I could only estimate covariances for the restricted sample of stockholders which is not necessarily more informative than the diagonal matrix.

2.5 Quantitative Results

2.5.1 2nd Stage Parameters

Table 2.3 reports the estimated second stage parameters. I take the bequest parameters from Cooper and Zhu (2016) which results in $L = 0.128$ and $\delta = 0.73$. My results imply that singles discount the future more (larger β) and display a slightly higher risk aversion (larger γ) than couples. In contrast, the estimated stock market participation cost is considerably larger for couples than it is for singles.

In the case that uses the identify matrix to weigh its moments, the estimate for the per period stock market participation cost corresponds to an annual cost of \$615.5 for couples and of \$326.5 for singles.[14] The coefficient of relative risk aversion is in contrast more similar: While my estimates suggest $\gamma = 2.677$ for singles, I find $\gamma = 2.654$ for couples. These values (especially for the coefficient of risk aversion) are at the lower end of estimates introduced by previous papers of portfolio allocation that augment a

14 As this is work in progress, I focus for now on the case that uses the inverse variance matrix only.

normally distributed return of the risky asset by a tail event. Generally, portfolio choice models have difficulties in matching participation rates and equity shares without introducing risk preferences or stock market participation costs that seem unlikely high when compared to life-cycle models without a portfolio choice. Introducing negative skewness in the return of the risky asset helps to address this puzzle. Moreover, in the current set-up, marriage and divorce introduce a dimension of financial risk for agents that so far has been largely overlooked in the household finance literature. Therefore, my model is able to match equity shares of households with relatively low degrees of risk aversion and standard values for the participation costs. Fagereng et al. (2017) estimate an annual stock market participation cost of \$69 but also introduce quite a high degree of risk aversion with $\gamma = 11$. In contrast, Catherine (2019) estimates a CRRA coefficient of $\gamma = 8.2$ and an annual stock market participation cost of $\$1,010$.[15] The estimates for β are in line with previous literature. Cooper and Zhu (2016) estimate a discount factor of 0.869, Fagereng et al. (2017) of 0.77 and Catherine (2019) of 0.92. However, given that my coefficients for the relative risk aversion are well below all of their estimates, my estimates for β are comparably low. One reason is that I exclude housing wealth and target the life-cycle profile of financial wealth instead of net worth. Moreover, the possibility of divorce increases the precautionary savings motives for couples while at the same time generating high-asset single households (who got divorced) that are absent in models with only bachelor households.

Table 2.3 Estimated 2nd Stage Parameters

β_c	β_s	γ_c	γ_s	S_c^F	S_s^F	$\mathcal{L}$	W
0.81	0.881	2.52	2.551	\$600	\$325	20874.76	$\mathcal{I}$
0.7911	0.8793	2.654	2.677	\$615.5	\$326.5	459.9713	$\frac{1}{\mathcal{V}}$

2.5.2 Model Fit

Figures 2.6 and 2.7 contrast the life-cycle profiles of equity shares and asset accumulation from the data (see section 2.2) with those generated by the model. The model is able to replicate the targeted evolution of wealth (Figure 2.6) for couples and single men. Moreover, it matches very well the un-targeted asset accumulation of single women over their life-cycle. Figure 2.7 illustrates the fit for the equity

15 In his paper, Catherine (2019) addresses the trade-off that life-cycle models of portfolio choice either require a very high degree of risk aversion (typically in combination with a very low discount factor) or a very high stock market participation cost to match the data by introducing cyclical skewness in labor earnings. To make his results comparable to mine, the listed values refer to the case when he estimates his model without cyclical skewness.

share: While the fit is quite good for couple, the model over-predicts the equity share of single men early in the life-cycle. However, most importantly, the model is able to capture the gap in equity shares between single men and single women (and hence, matches the life-cycle profile of equity shares for single women) without introducing preference heterogeneity by gender while at the same time matching overall asset accumulation by household structure.

When splitting the equity share along the extensive margin (participation rate) and the intensive margin (conditional risky share), I find that the model does a good job at matching the life-cycle profiles of participation rates, especially for single households (Figure 2.8). Finally, Figure 2.9 shows the model fit for life-cycle profiles of the conditional risky share. In contrast to the data, the model predicts the conditional risky share to be declining in age. This difficulty of portfolio choice models to match the life-cycle profiles of conditional risky shares is common: Because labor income is uncorrelated to the asset return, it acts as a substitute for the safe asset. Therefore, a decreasing human-to-financial wealth ratio over the life-cycle translates into a declining optimal risky share as individuals age. Nevertheless, the model correctly matches the average levels of conditional risky shares for singles.

Figure 2.6 Model Fit of Asset Accumulation

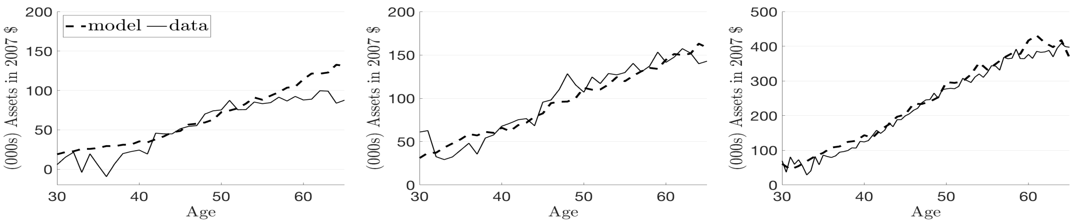

(a) Single Women **(b)** Single Men **(c)** Couples

Notes: Figure 2.6 plots the model fit of asset accumulation for single women, single men and couples. The solid lines show the data (as plotted in Figure 2.1) whereas the dashed line display the simulated life-cycle profiles generated from the model.

Figure 2.7 Model Fit of Equity Shares

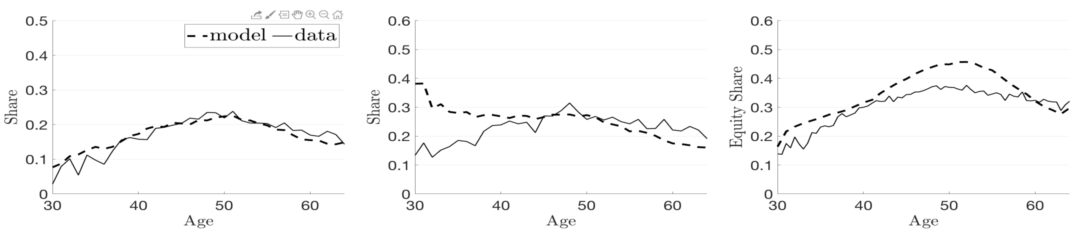

(a) Single Women **(b)** Single Men **(c)** Couples

Notes: Figure 2.7 plots the model fit of equity shares for single women, single men and couples. The solid lines show the data (as plotted in Figure 2.1) whereas the dashed line display the simulated life-cycle profiles generated from the model.

Figure 2.8 Model Fit of Participation Rates

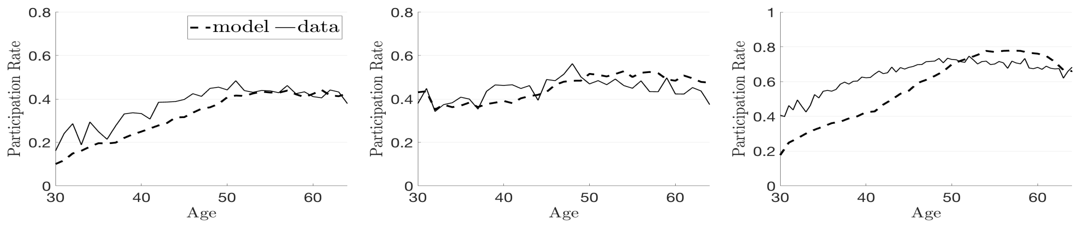

(a) Single Women **(b)** Single Men **(c)** Couples

Notes: Figure 2.8 plots the model fit of participation rates for single women, single men and couples. The solid lines show the data (as plotted in Figure 2.1) whereas the dashed line display the simulated life-cycle profiles generated from the model.

Figure 2.9 Model Fit of Conditional Risky Shares

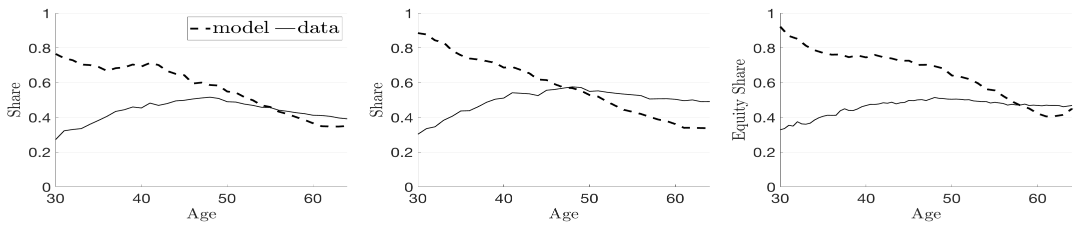

(a) Single Women **(b)** Single Men **(c)** Couples

Notes: Figure 2.9 plots the model fit of conditional risky shares for single women, single men and couples. The solid lines show the data (as plotted in Figure 2.1) whereas the dashed line display the simulated life-cycle profiles generated from the model.

2.5.3 Simulated Regressions

To compare the reduced-form results from Table 2.1 with those generated by the model, Table 2.4 replicates the same regression on simulated data. In particular, Column (1) shows the regressions estimates from model simulations whereas Column (2) re-reports the regression results from the main empirical specification (compare Table 2.1, Column (2)). All of these coefficients are un-targeted in the estimation exercise. The model slightly over-predicts the gender investment gap, especially early in life and hence, the coefficient for "single women" is more negative on the simulated data than it is on empirical data. In contrast, the interaction term of being a single woman and age is is more positive in the simulated dataset, resulting in an under-prediction of the investment gap as individuals age. When comparing the marginal effects, I find that the simulated data captures well the declining gender investment gap over the life-cycle. In particular, in both specifications, reduced form regressions that control for observable characteristics can explain gender differences in portfolio choices for individuals beyond age 45 but fail to fully explain the gap for younger households. Thus, the marginal effect for gender early in the life-cycle in reduced form regressions remains statistically significant if the underlying data generating process assumes preferences homogeneity across men and women. Hence, it appears that either factors which cannot as easily be controlled for (such risk exposure, expectations) or non-linearities explain the residual part of the gap. To uncover these factors to quantify their relative importance over the life-cycle, Section 2.6 performs several counterfactual exercises.

2.6 Counterfactual Simulations

2.6.1 Decomposing the Gender Investment Gap

In this Section, I decompose the gender gap in equity shares and in wealth levels along the dimensions of gender heterogeneity within the model, that is along income levels, income risk (productivity), marital transition probabilities, the expected characteristics of the partner in the event of marriage (the "marriage market": Π), the distribution across education levels, initial wealth levels at age 30, differences in household size (which is captured by the equivalence scale η) as well as medical expenses and age-dependent survival probabilities during retirement. In particular, I replace the female value for each channel with that of men and study the resulting gender gaps in asset holdings and in equity shares. Table 2.5 shows the results. The column "Model" reports the gender investment gap in the respective counterfactual scenario whereas the column "% explained" indicates how much of the baseline gap can be explained through the respective channel.

In general, aggregate portfolio allocations in the model are determined by the policy function for the optimal risky share $\alpha = \phi(X)$, conditional on state variables X, and the distribution of individuals across the state space. Thus, differences in investment behavior between the baseline model and the counterfactual scenario can arise because the distribution of individuals across the state space changes (*"composition effect"*) or because individual decision rules at any given point in the state space differ (*"policy effect"*).

Decomposing the Gap in Wealth Levels. Table 2.5 shows that differences in income levels, income risk and in household size explain the largest fraction of the wealth gap between single men and single women. Lower income levels naturally translate into less asset holdings, explaining 22.92% of the "gender wealth gap". At the same time, the income process of single women is less risky than that of single men. Therefore, assigning single women the male income risk increases female precautionary savings, reducing the gender gap in asset holdings. This channel in isolation explains on average 6.49% of the gap. Moreover, larger household sizes of single women act as a consumption commitment and lower the ability to save. On average, differences in household size between single men and single women explain 31.55% of the gender gap in wealth levels. Furthermore, giving single women the male marriage probabilities (that is, increasing their marriage hazard conditional on age) decreases the wealth gap by 9.41%. Especially during young age, agents in the model prefer marriage over being single because couples can pool their income and enjoy economies of scale for consumption. Increasing the likelihood of such a positive financial outcome consequently reduces precautionary savings.

The remaining channels are quantitatively less important for explaining gender heterogeneity in asset holdings. Assigning women the male medical expenses, the male survival probability or the male partner's characteristics in the event of marriage (marriage market) lowers asset holdings of single women and consequently increases the gender wealth gap. In the counterfactual scenario, lower medical expenses in very old age (beyond age 76) combined with a smaller chance of surviving up to that point decrease the incentive of single women to accumulate precautionary savings against longevity risk. However, as most gender differences in survival risk and in medical expenditures materialize at the very end of the life-cycle, this effect is quantitatively small when averaging over the working life.

Finally, when simulating the model under the assumption that both single men and single women start from the same (male) wealth level at age 30 substantially reduces the wealth gap early in life, increasing consumption of young single women.

Decomposing the Gap in Equity Shares. After having examined the gender gap in wealth levels, the next step is to decompose the gap in equity shares, that is the "gender investment gap". Naturally, as portfolio allocations are closely related to asset holdings, differences in both income levels and in household size not only explain the largest share of the gender gap in wealth levels but also in equity shares.

When single women receive the male income level, the simulated sample is composed of richer individuals who are more likely to cross the participation threshold of risky asset holdings. Figure 2.10a plots the policy function for the risky share from the baseline model (black line) and from the income counterfactual (red line). It shows that in addition to this *composition effect*, single women are also more willing to invest in the risky asset conditional on their wealth level (and other states). Because of the bond-like nature of labor income, a higher human capital endowment (i.e. more expected income in future periods) increases the willingness of single women to invest in the risky asset for a given level of wealth and current income. Thus, the effect of changing income levels on the gender investment gap operates both through the policy effect as well as through the composition effect. This result illustrates why controlling for current income in reduced-form regressions is not sufficient to explain the overall effect of income on portfolio choices. In addition to increased wealth through higher income today (and in previous periods), current portfolio choices are also affected by *expectations* over future income realization. Hence, it is rational for a single woman who has the same current income, wealth and education level as a single man to invest less risky because she expects a lower income in the future.

The same mechanism applies when lowering female household sizes to that of single men. The sample composition changes because smaller household sizes decrease per-period consumption and consequently translate into higher wealth levels. Moreover, the decision rules for a given point in the state space become more risky (see Figure 2.10b). Assigning single women the male household size not only decreases their consumption needs today but also in future periods, making them less vulnerable to financial risk and thus increasing their willingness to invest in the risky asset. Quantitatively, assigning single women the male income level reduces the gender gap in equity shares by 39.11%, whereas eliminating heterogeneity in household sizes narrows the gap by 41.16%.

In contrast, when single women face the same income risk as single men, the gender gap in equity share widens by 1.35%. The income process for single men is more volatile than for single women (see Table 2.2). Hence, giving single women the male income process lowers their willingness to invest in the risky asset. This "negative" effect on the gender gap in equity shares prevails despite the simulated sample being composed of on average richer single women (who are more likely cross the participation threshold) because of increased precautionary savings.

Moreover, if single women had the same marriage probability as single men (that is, increasing the likelihood of marriage), the gap in equity shares widens by 32.37%. The effect of increased marriage probability on equity shares mainly occurs because women hold less precautionary savings in the counterfactual scenario and are thus less likely to cross the threshold of risky asset participation. Moreover, especially early in life, it its optimal for single women to allocate a larger share of their portfolio towards the risky asset (conditional on participating) because they expect a higher household income (through marriage) in future periods. Generally, the significance of expected marriage for equity shares highlights the importance of considering marriage and divorce as a substantial financial risk when explaining the life-cycle behavior of portfolio choices. In particular, because marriage probabilities strongly differ by gender, they are key in explaining investment differences between single men and single women.

Figure 2.10 Policy Function for Risky Share (Single Women)

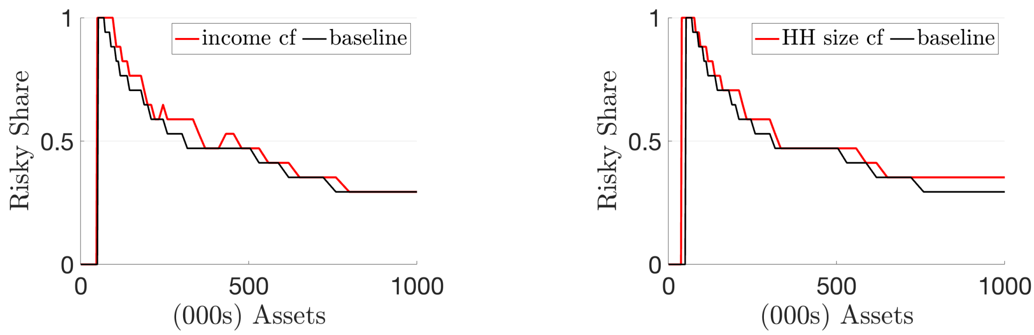

(a) Income counterfactual

(b) HH size counterfactual

Notes: Figure 2.10a plots the policy function for the risky share against current assets. The black line shows the policy function of the baseline model whereas the red line displays the policy function for the counterfactual scenario in which I assign single women the deterministic part of income process of single men ("income counterfactual") or the male household sizes ("HH size counterfactual"). The policy functions refer to single women with high education and medium-high labor productivity at age 40.

2.6.2 Composition vs. Policy Effect

Assigning single women the male income level increases the aggregate equity share of single women not only because their distribution across the spate space changes but also because it is optimal for them to invest more risky, conditional on state variables (see Figure 2.10a). Whereas I did control for the former when running reduced form regressions on empirical and simulated data (Table 2.4), it is not as straightforward to include expectations about future income levels in these regressions. The objective of this section is to quantify the relative importance of expectations versus current income differences on the gender investment gap along the life-cycle. To do so, I simulate the income counterfactual and restrict the policy functions of the risky share to be the same as in the baseline model. Consequently, any difference in life-cycle profiles between this simulation and the counterfactual with unrestricted policy functions can be attributed to the policy effect, that is, because single women choose a more risky portfolio allocation conditional on their state vector (through different expectations about future income). Figure 2.11 plots the results of this exercise. Figure 2.11a contrasts the life-cycle profiles of the baseline model (black solid line), the counterfactual with unrestricted policy functions (black dashed line) and the counterfactual in which I restrict the policy function for the risky share to be the same as in the baseline (red line). Any difference between the red line and the black dashed line can be attributed to the policy effect, thus, to changes in portfolio choices arising from different expectations. To quantify the relative importance of the policy effect vs. the composition effect, Figure 2.11b plots what percentage of the change in female equity shares in the income counterfactual can be explained by the composition effect (red line) and how much by the policy effect (black line).

I find that the policy effect almost entirely explains differences in equity shares between the baseline model and the income counterfactual early in the life-cycle. That is, larger equity shares for relatively young individuals arise from differences in decision rules (because of larger human capital endowments) rather than from differences in the sample composition. However, as individuals age, the composition effect becomes more important, eventually overtaking the policy effect at around age 56. Over the life-cycle, the remaining human capital endowment decreases and hence, its impact on policy functions becomes smaller. In contrast, higher income levels in previous years have translated into more savings, affecting the sample composition. In line with this finding, reduced form regressions (see Table 2.4) can explain the gender investment gap later in life whereas they fail do to so for younger households.

Figure 2.11 Composition vs. Policy Effect of Income Counterfactual (Single Women)

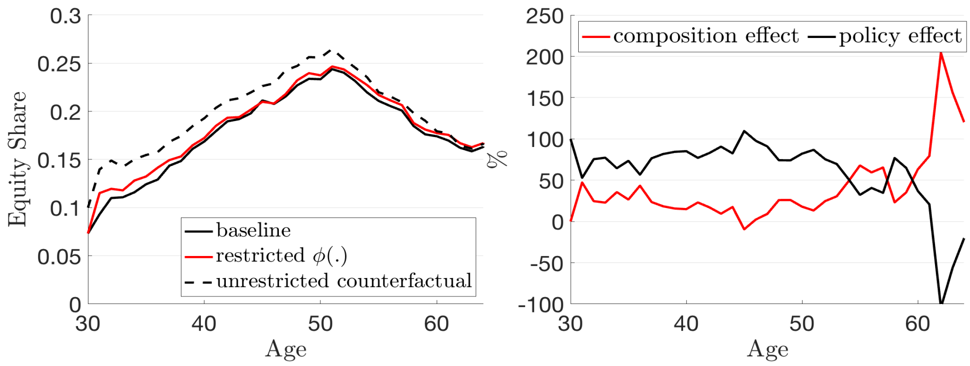

(a) Life-Cycle Profiles **(b)** Composition vs. Policy

Notes: Figure 2.11 decomposes the difference in female equity shares between the baseline model and the income counterfactual into a composition and into a policy effect. Figure 2.11a contrasts the life-cycle profiles from the baseline model (black solid line) to the unrestricted income counterfactual (black dashed line) and to the counterfactual that restricts the policy function for the risky share to be the same as in the baseline model (red line). In Figure 2.11b, the composition effect (red line) shows which percentage can be explained through on average richer individuals in the simulated sample whereas the policy effect (black line) shows which percentage can be explained by differences in decision rules for equity shares, fixing all other state variables; Both lines mechanically add up to 100 at every age.

Table 2.4 Regression Coefficients & Marginal Effects – Equity Shares of Singles

Equity Share	(1) Model Simulations	(2) Data SCF
single woman	-0.319***	-0.150***
	(0.0798)	(0.0166)
single woman*age	0.00363**	0.00159***
	(0.00157)	(0.000307)
age	-0.146**	0.0118
	(0.0646)	(0.0440)
$age^2 * 100$	0.383***	0.0134
	(0.138)	(0.0940)
$age^3 * 10000$	-0.298***	-0.0312
	(0.0965)	(0.0646)
High education	-0.0124	0.368***
	(0.0167)	(0.00560)
No. of HH members		-0.0511***
		(0.00293)
Income	0.472***	0.0422***
	(0.0112)	(0.00118)
Constant	-3.500***	-1.331**
	(0.984)	(0.646)
Observations	4,737	4,735
Year FE	No	Yes
ME for women at age 30	-0.101***	-0.0539***
	(0.0222)	(0.00380)
ME for women at mean age	0.0328	0.0113
	(0.0753)	(0.0150)
ME for women at age 65	0.154	0.0578**
	(0.127)	(0.0238)

Notes: Estimations are based on Tobit regressions on the sample of individuals that live in households with no spouse present. Column (1) are model simulations, column (2) refers to data from the SCF waves 1989 until 2016; individuals born between 1945 and 1960. Equity Share = Unconditional risky share. *single woman* is a dummy indicating that the household head is a women. *high education* is a dummy equal to one if the household head has more than 12 years of education. *safe assets* refers to safe liquid assets. "ME" indicates the marginal effect of being a women at the respective age. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table 2.5 Decomposition Results

Gap in (000s) Asset Holdings in 2007 $	Model	% explained	Data
Baseline	31.89		43.52
Male income level	24.58	22.92%	
Male income risk	29.82	6.49%	
Male HH size	21.83	31.55%	
Male marriage probability	34.89	-9.41%	
Male marriage market	33.33	-4.52%	
Male education distribution	32.78	-2.79%	
Male medical expenses	32.35	-1.44%	
Male survival probability	32.57	-2.13%	
Male initial wealth	26.04	18.36%	
Gap in Equity Share	**Model**	**% explained**	**Data**
Baseline	5.19%		5.93%
Male income level	3.16%	39.11%	
Male income risk	5.26%	-1.35%	
Male HH size	2.95%	43.16%	
Male marriage probability	6.87%	-32.37%	
Male marriage market	5.39%	-3.85%	
Male education distribution	5.47%	-5.39%	
Male medical expenses	5.05%	2.7%	
Male survival probability	5.45%	-5.0%	
Male initial wealth	4.18%	18.82%	

2.7 Conclusion

This paper studies the determinants of the gender investment gap over the life-cycle. It first provides empirical evidence that women allocate a smaller share of their liquid portfolio into risky assets while at the same time being less likely to hold any risky assets at all. Reduced form regressions reveal that the gender investment gap remains statistically significant after controlling for observable characteristics such as household size or income level, especially for young households. In contrast, an estimated structural portfolio choice model that restricts preferences to be equal across gender but allows for heterogeneity in observable characteristics and stochastic processes is able to (over-)explain the gap. Counterfactual simulations reveal that higher income levels of single men account for 39.11% in the observed gender gap in equity shares whereas gender differences in household sizes of singles explain 43.16%. Most importantly, the structural analysis finds that both contemporaneous income levels and household sizes as well as the *expected path* of these variables matter for current-period investment behavior. Because of the bond-like nature of labor income, a higher human capital endowment increases an agent's optimal equity share for any given level of wealth. Similarly, lower expected household sizes reduce future consumption needs and increase financial risk-taking. The effect of future realizations on portfolio choices is stronger for young households and hence, reduced form regressions that do not take into account households' expectations have troubles explaining the gender investment gap early in life.

Chapter 3

Joint Search over the Life-Cycle

Joint with Philipp Grübener and Lukas Nord

Abstract This paper studies how the added worker effect - intra-household insurance through increased spousal labor market participation - varies over the life cycle. We show in U.S. data that the added worker effect is much stronger for young than for old households. A stochastic life cycle model of two-member households with job search in a frictional labor market is capable of replicating this finding. The model suggests that a lower added worker effect for the old is driven primarily by better insurance through asset holdings. Human capital differences between employed young and old contribute to the difference but are quantitatively less important, while differences in job arrival rates play a limited role.

3.1 Introduction

Household earnings dynamics vary strongly over the life cycle. Recent literature documents that key moments of the earnings growth distribution exhibit significant age-dependency (De Nardi et al., 2019; Guvenen et al., 2021). Earnings variability is highest for young individuals as they change jobs frequently before settling into a stable job. However, the earnings growth distribution is more left-skewed for older individuals: Most of the time older individuals are employed in stable employment relationships at relatively high wages. If they lose this job, however, this fall off the job ladder implies very large earnings losses. In this paper we take a complementary perspective: Instead of investigating how risks change over the life cycle, we study how insurance against individual earnings risk varies over the life cycle. Specifically, we focus on an insurance margin against individual earnings and unemployment risk available to couples, the added worker effect (AWE), where a previously non-participating spouse enters the labor force upon job loss of the primary earner to stabilize joint earnings.

While the added worker effect has in general been widely documented, our focus on how it varies over the life cycle is novel to the literature.[1] Age differentials in the AWE are important for a variety of reasons: Observed heterogeneity along this margin improves our understanding of how well households at different ages are insured against income losses. Therefore, disparities in the availability of this self-insurance margin can alter the optimal provision of public insurance over the life cycle. Moreover, in light of demographic change any difference in the labor market behavior of old versus young households can alter aggregate labor market dynamics in the future.

We begin by providing empirical evidence on the added worker effect over the life cycle: Using data for the United States from the Current Population Survey (CPS), we show that the likelihood of a non participating spouse entering the labor force increases substantially when the primary earner loses her job compared to when she remains employed. We find, however, a strong age-dependency in this effect. In particular, the added worker effect is largest for young households and continuously declines over the life cycle. For the age group just before retirement, the added worker effect is almost non-existent. For young households, job loss of the primary earner is associated with a significant increase in the likelihood of an out of the labor force spouse entering the labor force both directly to employment and to unemployment. This finding is robust across education levels, the presence of children in the household, different reasons for being out of the labor force, different reasons for an employment to unemployment transition of the primary earner, and holds also when looking at only one cohort.

Still, there remain several candidate explanations for the observed change in the AWE over the life cycle. It might be that older households have accumulated sufficient asset holdings that allow them to smooth consumption during a potentially temporary job loss of the primary earner. In this case, older households do *not need* the added worker effect as an (additional) insurance margin. An out of the labor force spouse could in principle join the labor force, find employment, and stabilize joint earnings, but chooses not to do it. Alternatively, it could be that older spouses have been out of the labor force for a long time such that their labor market qualifications have become less valuable than those of younger individuals. In this case, spousal labor supply is *unavailable* as an insurance margin if the spouse can provide little marketable skills. In order to distinguish between the *need for* and the *availability of* the

[1] See the related literature below for a detailed discussion.

spousal insurance margin, we build a quantitative model of joint labor supply over the life cycle in a frictional labor market.

In the model, a household consists of two members, each of whom can be either employed, unemployed (and actively searching for a job), or out of the labor force. The labor market is frictional, an individual can only take up employment if she has a job offer. While both out of the labor force and unemployed individuals can receive job offers, unemployed members increase the chance of finding a job through costly search. Employed individuals face the risk of (exogenous) separation and wage changes due to match quality shocks. Human capital is accumulated while employed but depreciates during non-employment. A couple can jointly save in a risk-free bond. Job arrival rates are endogenous and determined by the solution to the vacancy posting problem of single-worker firms.

These model ingredients allow us to differentiate between the different candidate explanations for the age dependency in the added worker effect. Household savings are a key alternative insurance mechanism against individual unemployment risk. With a realistic life cycle savings profile the model can speak to whether differences in asset holdings between young and old are sufficient to explain the difference in the observed AWE. On the other hand, human capital accumulation and endogenous arrival rates allow for the possibility that older households might have fewer opportunities to provide insurance against individual risk, as human capital depreciates over long spells out of the labor force. Furthermore, firms might be less willing to hire older individuals as there is only little time remaining to recover hiring costs before their entry into retirement.

We calibrate the model to match key features of the U.S. labor market and of inequality over the life cycle. For the labor market, we focus on matching average transition rates across labor market states as well as the joint distribution of couples across labor market states. For inequality, we match life cycle income profiles and asset holdings over the life cycle. Without targeting them, the model reproduces reasonably well life cycle profiles of labor market transitions as well as very closely the age-dependency in the added worker effect. The model captures very well that the effect is largest for the young and smallest for the age group just before retirement.

With the calibrated model at hand, we perform counterfactuals to evaluate which mechanisms are important in explaining the age-dependency in the added worker effect. Our results suggest a significant influence of larger asset holdings of older households, which can serve as a cushion against temporary job loss. Higher human capital levels of old employed spouses relative to their younger counterparts – accumulated during a longer working life – make spousal labor supply less valuable as an insurance margin but are quantitatively less important. Differences in job arrival rates for young and old out of the labor force spouses play a limited role, as they turn out to be relatively low for both age groups.

In future work, we will evaluate the consequences of these mechanisms for the provision of optimal life cycle unemployment insurance. For such an analysis it is key to match the risk exposure of households over their life cycle as well as the private insurance mechanisms, which could be crowded out through public transfer payments. As our model covers a wide range of insurance mechanisms available to households at different stages of their life cycle, the framework naturally lends itself to this question. Michelacci and Ruffo (2015) study optimal life cycle unemployment insurance using a single earner life cycle search model.[2] They argue that unemployment insurance should be more generous for the young than for the old, as the insurance value is very high for individuals with little assets and the moral hazard problem is limited, as young individuals need to accumulate labor market experience. Studying this question in a search model of couples is relevant because unemployment insurance could crowd out the added worker effect, which is an important insurance margin among young households.

Related Literature. The added worker effect is widely studied in the empirical literature, going back to the seminal contribution of Lundberg (1985). The early literature following this paper does not find much evidence supporting the presence of the added worker effect in the data (Maloney, 1987, 1991). More recent literature, however, documents a positive added worker effect as a relevant insurance mechanism against the primary earner's job loss (Bredtmann et al., 2018; Guner et al., 2020; Halla et al., 2020; Stephens, 2002), using data for a variety of countries. Mankart and Oikonomou (2016b) and Mankart et al. (2021) show that the added worker effect has become more important in the U.S. over the last decades. The literature argues that the size of the added worker effect crucially depends on the institutional environment and the state of the business cycle. For example, Cullen and Gruber (2000) show that generous unemployment insurance crowds out a spousal labor supply response. Expanding upon previous work, we argue that there is a sizeable age-dependency in the added worker effect.

While the added worker effect has been studied extensively in the empirical literature, the vast majority of the large macro-labor literature focuses on the job search problem of a single earner household. Guler et al. (2012) is among the first papers to study the joint search problem of a couple by extending the classic single-agent search problems of McCall (1970), Mortensen (1970), and Burdett (1978). A number of recent papers introduces asset accumulation into the joint search framework, expanding on

2 Optimal age-dependent policies are also commonly studied in public finance. See for example Erosa and Gervais (2002), Weinzierl (2011), and Heathcote et al. (2020).

the single agent search problem with asset accumulation as in Lentz (2009), Krusell et al. (2010), and Krusell et al. (2017). The focus of these papers is mostly on business cycle dynamics. Mankart and Oikonomou (2016a) build a search model with two member households to explain the cyclical properties of employment and labor force participation. Wang (2019) builds a model showing that joint household search is crucial for accounting for the countercyclicality of womens' unemployment rate. Ellieroth (2019) argues that there is precautionary labor supply by spouses whose partners face an increased job loss risk in recessions. Garcia-Perez and Rendon (2020) focus on the role of household wealth for the added worker effect. Birinci (2019), Choi and Valladares-Esteban (2020), and Fernández-Blanco (2020) investigate the implications of joint search for optimal unemployment insurance. Bardóczy (2020) focuses on the role of spousal labor supply as an automatic stabilizer for aggregate consumption. Relative to these papers, we focus on the life cycle dimension of the joint search problem to analyze whether the age-dependency in the added worker effect is explained by changing opportunities or changing insurance margins.

Life cycle search problems have been studied in the literature, but mostly in single earner frameworks. Chéron et al. (2011, 2013) extend the random search framework of Mortensen and Pissarides (1994) to a life cycle setting. Menzio et al. (2016) build a directed search life cycle model in the tradition of Moen (1997) and Menzio and Shi (2011). Griffy (2021) extends their model by incorporating risk averse workers and borrowing constraints. More closely related to our paper, Haan and Prowse (2017) propose a structural life cycle model of labor supply, consumption, and savings of married couples. They focus on the optimal mix of unemployment insurance and social assistance but do not discuss any age-dependency in the added worker effect. Finally, the current paper is related to a number of studies analyzing life cycle labor supply decisions of couples in incomplete market frameworks (Blundell et al., 2016; Ortigueira and Siassi, 2013; Wu and Krueger, 2021).

Roadmap. The paper proceeds as follows. Section 3.2 contains the empirical evidence. In Section 3.3 we introduce the model setup. Section 3.4 contains the calibration and section 3.5 the results. Section 3.6 concludes.

3.2 Evidence

We begin by providing evidence on the added worker effect from U.S. micro data. The following section first explains the data and the sample selection criteria. In a next step, we provide empirical evidence of the AWE in our sample and show that its magnitude is decreasing in age.

3.2.1 The Sample

To compute joint labor market transitions, we work with data from the Current Population Survey (CPS), provided by the Integrated Public Use Microdata Series (IPUMS) (Flood et al., 2020).[3] The CPS is a monthly rotating panel which is representative for the U.S. population. Households enter the survey for four consecutive months, drop out for eight months, and are re-interviewed for another four months. In our setting, the unit of observation is a couple. Our final sample spans from 1994 until 2020 (pre-Covid) and is restricted to couples who are both between 25 and 65 years old. We mainly focus on couples with one spouse working and the other spouse out of the labor force. We include both legally married as well as cohabiting couples, irrespectively of their sex. In contrast, we drop couples who report that one spouse lives permanently outside of the household or is institutionalized. Moreover, we only keep couples for whom we observe the labor market status of both spouses in every month that they are interviewed. Throughout the analysis, we weigh each observation by the provided survey weights.

3.2.2 Uncovering the AWE from Joint Labor Market Transitions

We follow Guner et al. (2020) in our method to calculate the added worker effect from the data. First, we classify all individuals either as *employed* (E), *unemployed* (U) or *non-participating* (N) as outlined in the CPS. Hence, there exist nine possible combinations of labor market states for each couple. A common issue when considering multiple non-employment states is misclassification between unemployment and non-participation, resulting in implausibly high transition rates across these two. We therefore adjust labor market flows as in Elsby et al. (2015) and re-classify individuals who report to be unemployed (non-participating) in one month but to be out of the labor force (unemployed) in both the following and in the previous month as non-participating (unemployed).

[3] In addition, this paper uses data from the Panel Study of Income Dynamics (Panel Study of Income Dynamics, 2021).

In a next step, we pool all observations and construct a 3×3 matrix of joint labor market transition probabilities, conditional on the couple having one member previously employed and one out of the labor force. Table 3.1 and Table 3.3 display our main results. In each table, the columns refer to the monthly labor market transition of the household's primary earner, that is either employment-to-employment (EE), employment-to-unemployment (EU), or employment-to-non-participating (EN). In contrast, each row indicates the probability of the spousal labor market transition, conditional on the respective transition of the primary earner. Given that for this exercise we only include couples with one member employed and the other one non-participating, spouses can either transition from non-participating to employment (NE), from non-participating to unemployment (NU) or remain out of the labor force (NN). We define the added worker effect as the change in the conditional probability of the spouse transitioning from non-participating to employment (NE) or from non-participating to unemployment (NU) if the primary earner becomes unemployed (EU) in contrast to when the primary earner remains employed (EE). Referring to Table 3.1, we compute the added worker effect as the difference between the second and first column, adding up the first and the second row.

Overall Effect

Table 3.1 shows the overall strength of the added worker effect in our sample. The likelihood that a spouse enters the labor force increases by 5.9 percentage points, if the primary earner becomes unemployed compared to when the primary earner remains employed, confirming the existence of the added worker effect in our sample.[4] This result is in line with Guner et al. (2020), who find an overall AWE of 6.89 percentage points with CPS data spanning from 1976 to 2018 for couples between 25 and 54 years.

Zooming in on the precise margin of adjustment, we find that the conditional probability of the spouse transitioning directly into employment increases by 1.98 percentage points, whereas the conditional probability of the spouse transitioning into unemployment increases by 3.92 points. Thus, around two thirds of the overall AWE arise from individuals transitioning into unemployment, highlighting the importance of explicitly distinguishing between unemployed and non-participating individuals.

[4] In this paper we focus on the transitions of out of the labor force spouses conditional on the labor market transitions of primary earners. In the appendix, Tables C.1 and C.2 we also report the conditional transition probabilities of unemployed and employed spouses, respectively. There is a slightly higher likelihood that unemployed spouses transition to employment or stay unemployed rather than leave the labor force if the primary earner loses the job compared to the primary earner staying employed. However, evidence for insurance through spousal labor supply is strongest when considering out of the labor force spouses, which we focus on.

Table 3.1 Joint Labor Market Transitions (Full Sample)

	Primary earner transition		
	EE	EU	EN
Cond. prob. of spousal NE transition	6.03%	8.01%	16.79%
Cond. prob. of spousal NU transition	1.63%	5.55%	1.33%
Cond. prob. of spousal NN transition	92.34%	86.44%	81.88%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions for the entire population.

Table 3.2 AWE by reasons of Unemployment for Household Head

	EE	EU (by reasons for U)			
		Layoff	Job Loser	Temp. Job ended	Job Leaver
NE	6.03%	6.13%	8.81%	7.56%	10.47%
NU	1.63%	3.51%	6.66%	6.59%	7.68%
NN	92.34%	90.35%	84.53%	85.85%	81.86 %

Notes: This table shows the added worker effect (as defined in the main text) by reason for the EU transition of the primary earner.

Some couples may wish to leverage spousal labor supply as an insurance margin against job loss but labor market frictions (or the lack of appropriate job offers) prevent them from doing so. If we only considered transitions from non-employment into employment, we would hence understate spousal labor supply adjustments in response to the job loss of the primary earner.

To further investigate the added worker effect, Table 3.2 splits primary earners by the reason for why they became unemployed. In particular, we distinguish between laid-off workers (who face a high chance of being recalled), job losers, workers whose temporary contracts ended, and voluntarily job leavers. Table 3.2, which splits the EU transition of the primary earner by reason for entering unemployment, shows that our finding is not solely driven by household members voluntarily quitting (column *Job Leavers*, especially with spouse NE) upon employment of their partner. The effect for those exogenously separated (*Job Losers*) is of similar magnitude, with a slightly decreased AWE for households in which the head's job loss can be seen as expected (*Temp. Job ended*) or as temporary in nature (*Layoff*).

While in the main text we focus on couples where one spouse is employed and the other is out of the labor force, in Appendix C.1 we include similar tables for couples that start as both employed or with one employed and one unemployed member. We can also see in these transition matrices that unemployed spouses are slightly more likely to enter employment or keep looking for jobs rather than dropping out of the labor force if the primary earner moves from employment to unemployment compared to when the primary earner stays employed. However, the main pattern that emerges from

Table 3.3 Joint Labor Market Transitions by Age

	Primary earner transition		
	EE	EU	EN
Age Spouse 25-35:			
Cond. prob. of spousal NE transition	6.66%	9.30%	26.93%
Cond. prob. of spousal NU transition	2.00%	6.89%	2.02%
Cond. prob. of spousal NN transition	91.34%	83.81%	71.05%
Age Spouse 36-45:			
Cond. prob. of spousal NE transition	6.73%	9.32%	26.69%
Cond. prob. of spousal NU transition	1.86%	6.37%	2.00%
Cond. prob. of spousal NN transition	91.41%	84.31%	71.30%
Age Spouse 46-55:			
Cond. prob. of spousal NE transition	6.13%	7.96%	16.62%
Cond. prob. of spousal NU transition	1.62%	4.79%	1.72%
Cond. prob. of spousal NN transition	92.25%	87.25%	81.66%
Age Spouse 56-65:			
Cond. prob. of spousal NE transition	4.29%	3.73%	8.69%
Cond. prob. of spousal NU transition	0.90%	2.75%	0.56%
Cond. prob. of spousal NN transition	94.81%	93.52%	90.76%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions by age group.

these two transition matrices is that couples often make joint transitions: The likelihood of a spouse dropping out of the labor force is drastically increased when the primary earner also transitions from employment or unemployment to out of the labor force.

The Added Worker Effect by Age

To analyze the life cycle dimension of the added worker effect, we split our sample into four age brackets and construct joint labor market transitions for each group in the same manner as above. Table 3.3 displays the results. We find a strong age-dependency in the strength of the AWE: For the youngest group (25 to 35 years), the likelihood that the spouse enters the labor force upon the job loss of the primary earner increases by 7.53 percentage points, for the young middle aged (36 to 45 years) it increases by 7.10 points, for the older middle aged (46 to 55 years) by 5.00 points, and eventually only slightly increases by 1.29 points for the oldest group (56 to 65 years). Thus, spousal labor supply adjustments of the youngest age group are more than five times larger than for the oldest age group.

For the young, we find behavioral responses both from non-participating directly into employment (2.64 percentage points) as well as into unemployment (4.89 percentage points). Thus, the relative share of young individuals transitioning directly into employment is slightly larger than for the entire sample. In contrast, for the oldest age group, we only find small behavioral responses into unemployment (1.85 percentage points) and no response directly into employment (-0.56 points).

3.2.3 Dynamic Response

So far, we have focused on the contemporaneous spousal labor supply response, that is, the probability that a spouse enters the labor force in the *same month* as the head transitions into unemployment. This most likely understates the overall strength of the added worker effect since spousal labor supply responses may occur in prior months (anticipation effects) or with some delay. In fact, Ellieroth (2019) documents spousal insurance not only in response to actual job loss of the primary earner but also in anticipation of such event, a phenomenon that she names "precautionary labor supply". To analyze the strength of both anticipation and lagged responses, we run the following linear regression specification:

$$\Delta LFS_{it}^{sp} = \alpha_j + \beta_j \Delta ES_{it+j}^{h} + \gamma_j X_{it} + \epsilon_{jit}, \tag{3.1}$$

where ΔLFS_{it}^{sp} is a dummy that takes the value 1 if the non-participating spouse of couple i transitions either into employment or into unemployment between month $t-1$ and t, and 0 if she or he remains out of the labor force. Similarly, ΔES_{it}^{h} is defined as a dummy taking the value 1 if the primary earner transitions from employment into unemployment whereas it is 0 if the head stays in employment. X_{it} further controls for month fixed-effects, year fixed-effects, state fixed-effects, sex, race, education, children as well as the quarterly unemployment rate in the couple's state of residence.

Our coefficient of interest is β_j, indicating the likelihood that the spouse enters the labor force in month t if the household head transitions into unemployment in month $t+j$ versus when he or she remains employed (i.e. the strength of the AWE in month $t+j$). We conduct the analysis for $j = \{-2, -1, 0, 1, 2\}$. In the CPS, we observe the same couple for at most four consecutive months and hence a maximum of three consecutive labor market transitions, preventing us from considering more distant leads and lags. Figure 3.1 reports the results for the entire sample, whereas Figure 3.2 splits the observations by age.

In line with section 3.2.2, Figure 3.1 confirms the overall strength of the AWE of around 6.1 percentage points in the contemporaneous month. Moreover, this effect is statistically significantly different from zero. In addition to the contemporaneous effect, we find strong support of both anticipation and lagged effects, albeit of lower magnitude. Our results indicate that spousal labor supply responses in the months preceding and in the months after the primary earner's job loss are around half as strong as the direct response. When splitting the sample by age (Figure 3.2), we find that the contemporaneous effect is

Figure 3.1 Δ Pr(Spouse enters LF) <u>this month</u>

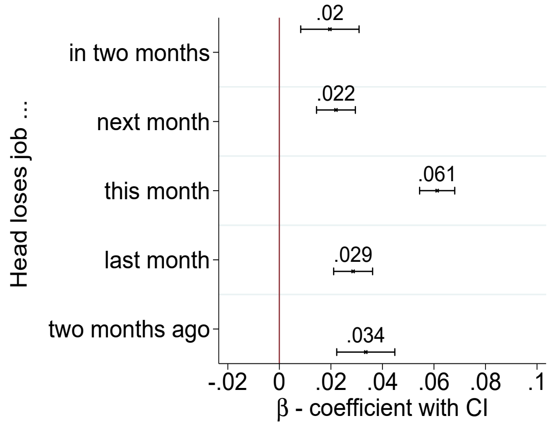

Notes: Figure 3.1 shows the change in probability that a non-participating spouse enters the labor force (either as unemployed or as employed) if the household head loses/lost the job in two months, next month, this month, last month or two months ago, respectively, relative to the baseline in which the household head remains employed. The sample includes couples in which one spouse is working and one spouse is out of the labor force between age 25 and 65 from the Current Population Survey (CPS), waves 1994 until 2020. The regression producing the coefficients is Equation 3.1.

statistically significant for all age groups, however it is around five times stronger for the young than for the old. Moreover, young households display both lagged responses as well as anticipation effects, whereas we cannot confirm any clear pattern of those among households between 56 and 65 years. We relegate the results for the two middle age groups to Figure C.1 in the appendix.

Lastly, in Figure 3.3, we again split the sample by reasons for unemployment of the primary earner (as in Table 3.2). Generally, the figure confirms that the probability that a non-participating spouse enters the labor force increases most if the EU transition of the primary earner is due to a quit or job loss, and less so in case of a layoff when there is a chance of being recalled. Interestingly, for spouses of household heads who voluntarily leave their job the effect two months ahead and the two month lagged effect are smaller, while the effect in the month before and after the primary earner transition is larger. This finding can be taken as indication that these labor market transitions are coordinated choices within a short time span.

3.2.4 Robustness

In this section, we explore further channels that could result in the observed age-dependency in the added worker effect without relating to life cycle heterogeneity in the insurance value of the AWE itself

Figure 3.2 Δ Pr(Spouse enters LF) <u>this month</u>

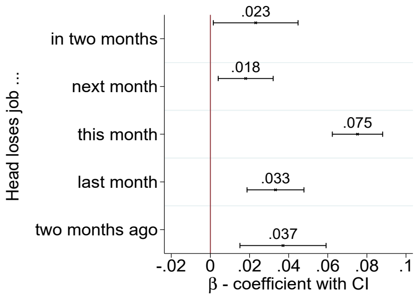

(a) Age 25 to 35

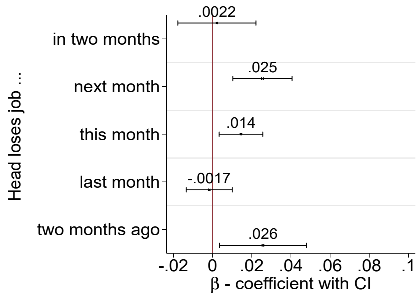

(b) Age 56 to 65

Notes: Figure 3.2 shows the change in probability that a non-participating spouse enters the labor force (either as unemployed or as employed) if the household head loses/lost the job in two months, next month, this month, last month or two months ago, respectively, relative to the baseline in which the household head remains employed. The sample includes couples in which one spouse is working and one spouse is out of the labor force between age 25 and 35 (Figure 3.2a) and between age 56 and 65 (Figure 3.2b) from the Current Population Survey (CPS), waves 1994 until 2020. Age refers to the non-participating spouse. The regression producing the coefficients is Equation 3.1.

Figure 3.3 Δ Pr(Spouse enters LF) <u>this month</u>

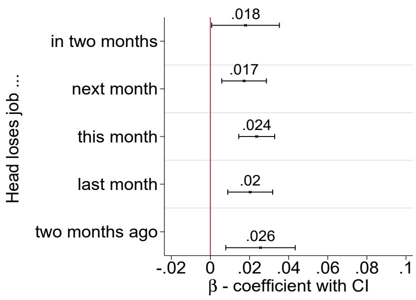

(a) Job Loser – On Layoff

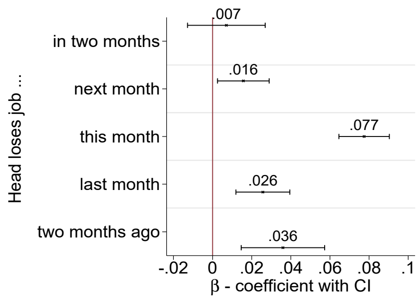

(b) Job Loser

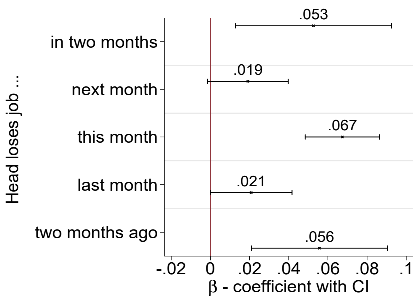

(c) Temporary Job Ended

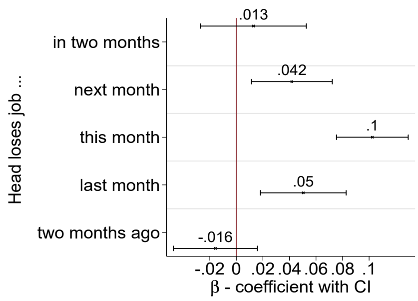

(d) Job Leaver

Notes: Figure 3.3 shows the change in probability that a non-participating spouse enters the labor force (either as unemployed or as employed) if the household head loses/lost the job in two months, next month, this month, last month or two months ago, respectively, relative to the baseline in which the household head remains employed; split by reasons for unemployment of the household head. Specifically, Figure 3.3a shows the results if the household head is on layoff, Figure 3.3b if the household head lost his job, Figure 3.3c if a temporary job ended and Figure 3.3d if the head voluntarily quit his or her job. The sample includes couples in which one spouse is working and one spouse is out of the labor force between age 25 and 65 from the Current Population Survey (CPS), waves 1994 until 2020. The regression producing the coefficients is Equation 3.1.

nor to other insurance margins that differ by age.[5] All corresponding tables are listed in Appendix C.1.

Education. If educational attainment differs by age and at the same time affects spousal labor supply responses, the stronger AWE for younger couples may simply arise from differences in education levels between old and young couples.[6] Indeed, Table C.3 confirms that the AWE is larger for spouses with a college degree. However, when splitting the sample by age and education (Panel III to VI in Table C.3), the decreasing magnitude of the AWE over the life cycle holds both among spouses with a college degree and among those without a college degree.

Cohort Effects. If preferences for labor supply or within household insurance differ by cohorts, any age-dependency in the added worker effect between old and young couples may be driven by these underlying preference shifts. Female labor force participation increased substantially between the 1960s and the 1990s. Hence, entering the labor force upon the head's job loss may be easier for young couples if deviations from the traditional family model are societally more accepted. We address this concern in two ways. First, we split our sample by gender and age. Male labor force participation changed to a much lesser extent than that of women. If we can replicate the age-dependency in the AWE for couples in which the non-participating spouse is a man, possible cohort effects are less concerning. Table C.4 (Panels I and II) shows the results of this exercise. Although we find that the overall probability of the spouse joining the labor force is higher when the non-participating household member is a man, we do not find significant changes in the strength of the AWE (i.e. in the increased likelihood that the spouse enters the labor force when the household head becomes unemployed, compared to when the head remains employed). Focusing only on male non-participating spouses, young households still show a stronger AWE than older couples. We take this as suggestive evidence that our results are not driven by changing patterns of female labor force participation.

Arguably, couples for which a man is non-participating could be a particular selection whose preferences differ from those of the remaining population. To address this concern, we extract one cohort and repeat the empirical exercise on this restricted sample. In particular, we focus on couples in which the non-participating spouse was born between 1960 and 1970. We choose this timespan to ensure sufficiently many observations both for the young and for the old age brackets. Table C.4 (Panel III and

[5] Some of these variables are also included as controls in the regressions. We still address the economically most important ones explicitly in this section.

[6] Generally, heterogeneity in education levels by age is low: around 45% of spouses among the youngest age group have a college degree, whereas around 40% of spouses among the oldest age group do.

IV) reports the results. Again, we can confirm the decreasing magnitude of the AWE over the life cycle for this particular cohort, i.e. for the same cohort when young and when old.

Children. Young couples are more likely to have children living in their household, which arguably affects labor supply behavior and could therefore result in the observed differences of spousal labor supply insurance. On the one hand, couples with children might have stronger incentives to enter the labor force in response to the job loss of the primary earner because they have larger consumption commitments and stronger saving motives (e.g. saving for college). On the other hand, if household members specialize in childcare and paid work, the willingness of the spouse who specializes in childcare to enter the labor force might be low. To address this issue, Table C.5 reports the AWE for couples below age 40 (to avoid picking up age-effects) with and without children as well as for couples below age 40 with and without children under age five (who require the most childcare). Our results indicate that out of the labor force spouses in couples without children have a higher baseline probability of entering the labor force, independently of the labor market transition of the primary earner. However, we do not find any (significant) differences in the overall strength of the AWE between couples with and without children across both specifications.

Reasons for Non-Participation. Individuals do not participate in the labor force for a variety of reasons that are age-dependent. At the same time, the reason for being out of the labor force can affect the strength of the added worker effect. For example, if the non-participating spouse is retired, transitioning back into the labor force has a much smaller insurance value because of pension payments. Similarly, if the non-participating spouse dropped out because of bad health, she or he might simply not be able to start working if the primary earner becomes unemployed. Arguably, both retirement and health related non-participation are more prevalent among the old. Therefore, Table C.6 repeats the empirical analysis excluding retired spouses (Panels I and II), disabled or ill spouses (Panels III and IV), as well as excluding both retired and disabled/ill spouses (Panels V and VI). Unsurprisingly, these restrictions do not impact our baseline results for the young age group. However, we also do not find any significant impact on the strength of the AWE among the old. If anything, spouses are more likely to join the labor force in general when excluding retirees, however, the increase in the likelihood of entering (un)employment in response to the primary earner's job loss is not larger (or smaller) when repeating the analysis on the three subsamples.

Business Cycle. As much of the joint search literature focuses on business cycle dynamics (e.g. Mankart and Oikonomou (2016a) and Birinci (2019)), we investigate in this exercise whether our results differ by the state of the economy. In Table C.7 we split the sample by NBER recessions and expansions. The state of the business cycle might matter for the added worker effect in several ways. On the one hand, if a primary earner loses a job in a recession, it might be harder to find a job again, so that insurance through spousal labor supply could be more important. On the other hand, it could also be harder for an out of the labor force spouse to find a job and provide this insurance. We do not, however, find large differences in the AWE across young and old for different states of the business cycle.

Income. A deficiency of the CPS for our analysis is that we do not observe asset holdings of households, which are another key insurance margin available to them. We have, however, some information on total income of a couple over the past year. Total income may proxy for the ability of households to build up savings, but it is also correlated with other characteristics such as education. We split couples into income terciles and compute transition matrices for these different income groups. Table C.8 reports our findings. Pooling all age groups we observe a sizeable AWE for low and high income groups. For the old, the added worker effect is relatively weak for both low and high income groups. When only considering the young, the AWE is smaller for the high income group than for the low income group. This results may reflect that among high income couples the primary earner has a higher chance of being reemployed or that the high income group has larger savings. Both these channels will be present in the our quantitative theory, to which we turn next.[7]

3.3 Model

The empirical evidence presented so far suggests that there is a significant age-dependency in the added worker effect: Spousal labour supply is a more important insurance margin for young than for old couples. We now build a life cycle search model with two-member households in order to better understand why the added worker effect is more prevalent among the young.

[7] In ongoing work, we extend our empirical analysis using data from the Survey of Income and Program Participation (SIPP). We started with the CPS as it is the main source for monthly labor market statistics in the United States. The SIPP, however, has the advantage that we can observe households' asset holdings.

Figure 3.4 Human Capital Transitions

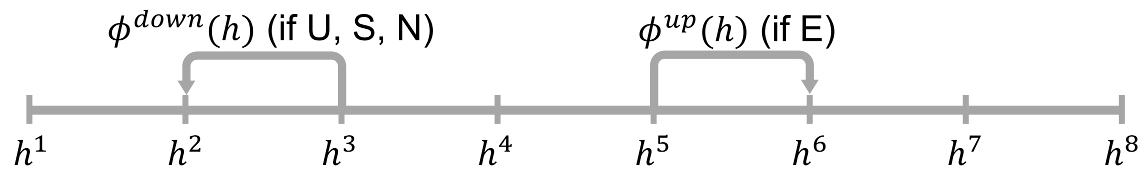

Notes: Figure 3.4 illustrates human capital transitions in the model.

3.3.1 Environment

The economy is populated by two-member households. We assume that both members have the same age. Households live for T periods, after which they die deterministically. Households retire jointly after a working life of T_W periods, so that retirement lasts $T - T_W$ periods.

During working life an individual can be in one of four labor market states. An individual can be employed (E), in which case the agent receives a wage payment. If the individual does not have a job, there are three other labor market states: First, an agent may be unemployed and receive benefits (U). Second, the agent can be unemployed without receiving benefits (S). In both these states, the agent exerts costly search effort in order to increase the probability of finding a job. Third, an agent may choose to not exert this costly search effort. In that case, the agent is considered to be out of the labor force (N). Individuals who are not actively searching can never receive unemployment benefits. Given these four individual labor market states, there are 16 joint labor market states for a two-member household: $jk \in \mathcal{J} = \{E,U,S,N\} \times \{E,U,S,N\}$.

Each household member is endowed with a level of human capital, which evolves stochastically depending on the agent's employment status and current human capital level. If an individual member is employed, the human capital will go up by one unit with probability $\phi^{up}(h)$. For non-employed agents, human capital drops by one unit with probability $\phi^{down}(h)$. This process is illustrated in Figure 3.4.

While employed, an individual is additionally characterized by match quality z, which evolves according to a first-order Markov process. The match quality and the human capital level jointly determine the wage an individual receives. Non-employed individuals do not have a match quality, however they draw one upon finding a new job.

Individual labor market transitions are illustrated in Figure 3.5. An employed agent can receive an exogenous separation shock with probability $\delta(h)$, which depends on the level of human capital. If such a separation shock occurs, the agent transitions to unemployment and receives unemployment benefits. Note that in case of a separation shock an agent can choose to immediately leave the labor force instead of becoming unemployed and receiving benefits. This can be beneficial because no costly search effort is exerted while out of the labor force. If there is no separation shock, the individual can choose between staying employed and quitting. If she chooses to quit, she can either become unemployed without receiving benefits or leave the labor force entirely.

An unemployed agent who receives benefits can transition to all other labor market states. First, she receives a job offer with probability $\lambda^U(x_i)$ and transitions to employment if she chooses to accept the offer. The arrival rates with which non-employed agents receive job offers are endogenously determined as the solution to an optimal vacancy posting problem of firms (see below) and for household member i depend on state $x_i = \{h_i, h_{-i}, z_{-i}, a', jk\}$. An agent can also choose to reject the offer and might do so if the initial match quality draw is low. In that case, it may be preferable to wait for a new offer with a potentially better match quality draw. Second, an unemployed worker who receives benefits can stochastically lose benefit eligibility with probability ϕ^{US}, capturing that unemployment benefits run out after a certain time period. Third, she can choose to stop searching and leave the labor force. Similarly, an unemployed worker without benefits receives job offers with probability $\lambda^S(x_i)$ and can quit the labor force.

Finally, out of the labor force agents receive job offers with probability $\lambda^N(x_i)$, even though they do not exert active search effort. This assumption is necessary to capture the empirical observation that individuals directly transition from out of the labor force into employment. Moreover, non-participating agents can rejoin the labor force as unemployed without benefits.

While each household member has an individual labor market state, human capital level, and match quality shock when employed, households jointly have access to a risk-free bond. They can save in this bond at the exogenous interest rate r. Borrowing is not allowed.

3.3.2 Household Search Problem

Timing in the model is as follows: In each period, households first receive their labor income (wages or unemployment benefits) as well as their asset income from investing in the risk-free bond. Given their budget constraint, households then make a consumption-savings choice. Afterwards, separation shocks, job offers and potential losses of benefit eligibility are realized for both household members in parallel. Next, match quality shocks and human capital transitions are revealed. Finally, households

Figure 3.5 Labor Market Transitions in the Model

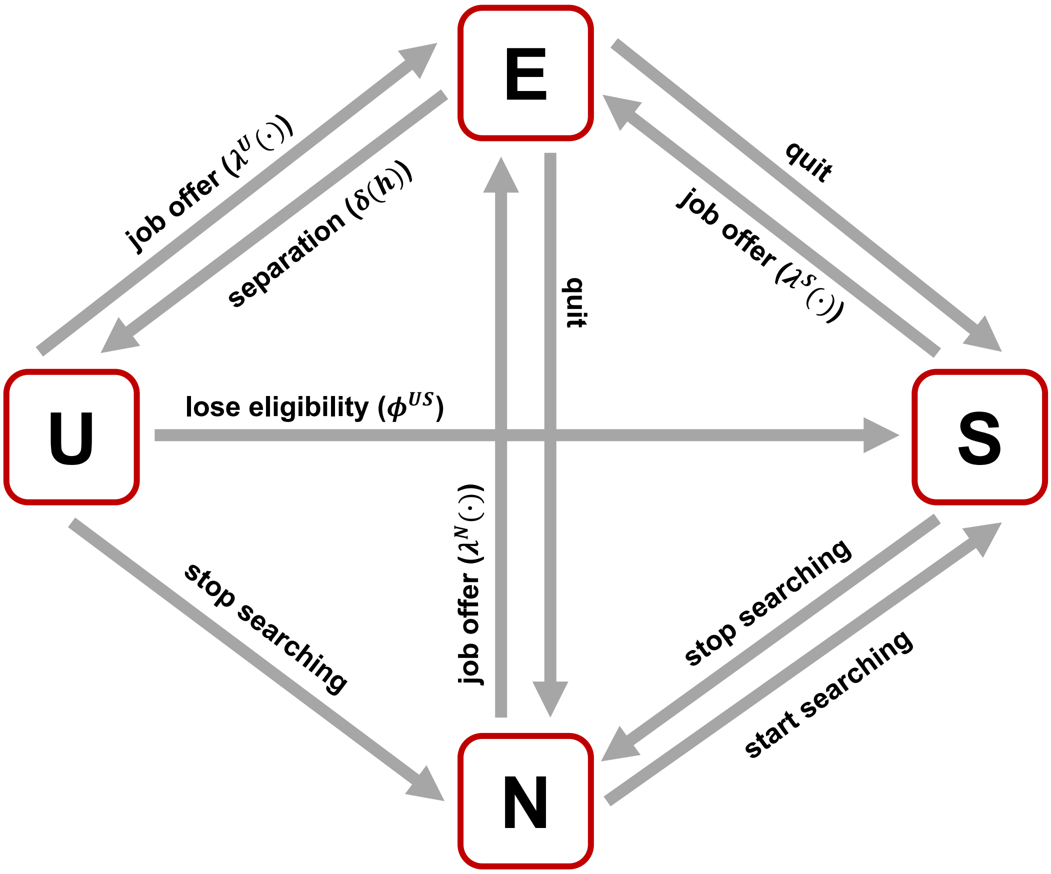

Notes: Figure 3.5 illustrates possible labor market transitions in the model. $x_i = \{h_i, h_{-i}, z_{-i}, a', jk\}$ is the relevant state for the arrival rate of household member i.

choose their joint future labor market state from the feasible subset of $\mathcal{J}$, which is determined by their previous labor market state and job offers, separations, and benefit eligibility losses.

Table 3.4 summarizes all possible combinations of job opportunities and unemployment benefit eligibility of the two household members along with the associated choice sets over joint labor market states. The superscripts to $\mathcal{J}$ indicate whether the household members have the opportunity to be employed. An employment opportunity arises either because an agent was employed in the previous period and did not receive a separation shock or because an agent received a job offer while non-employed. If both members have the opportunity to be employed, the superscript is EE. In contrast, X indicates that a member cannot be employed. Hence, EX and XE are the cases in which only one member has a job opportunity, whereas XX indicates that neither household member can be employed in the following period.

Table 3.4 Labor Supply Choice Sets

Benefit Eligibility	**Job (Offer)** Both	Member 1	Member 2	None
Both	$\mathcal{J}_{UU}^{EE} = \{E,U,N\} \times \{E,U,N\}$	$\mathcal{J}_{UU}^{EX} = \{E,U,N\} \times \{U,N\}$	$\mathcal{J}_{UU}^{XE} = \{U,N\} \times \{E,U,N\}$	$\mathcal{J}_{UU}^{XX} = \{U,N\} \times \{U,N\}$
Member 1	$\mathcal{J}_{UX}^{EE} = \{E,U,N\} \times \{E,S,N\}$	$\mathcal{J}_{UX}^{EX} = \{E,U,N\} \times \{S,N\}$	$\mathcal{J}_{UX}^{XE} = \{U,N\} \times \{E,S,N\}$	$\mathcal{J}_{UX}^{XX} = \{U,N\} \times \{S,N\}$
Member 2	$\mathcal{J}_{XU}^{EE} = \{E,S,N\} \times \{E,U,N\}$	$\mathcal{J}_{XU}^{EX} = \{E,S,N\} \times \{U,N\}$	$\mathcal{J}_{XU}^{XE} = \{S,N\} \times \{E,U,N\}$	$\mathcal{J}_{XU}^{XX} = \{S,N\} \times \{U,N\}$
None	$\mathcal{J}_{XX}^{EE} = \{E,S,N\} \times \{E,S,N\}$	$\mathcal{J}_{XX}^{EX} = \{E,S,N\} \times \{S,N\}$	$\mathcal{J}_{XX}^{XE} = \{S,N\} \times \{E,S,N\}$	$\mathcal{J}_{XX}^{XX} = \{S,N\} \times \{S,N\}$

Notes: This table shows the labor supply choice sets of households.

The logic for the subscripts is similar. However, they refer to unemployment benefit eligibility of the individual household member. Again, U indicates eligibility, while X refers to non-eligibility.

We are now in the position to formally state the household search problem. The value function of a household of age t in joint labor market state jk is

$$V_t^{jk}(z,h,a) = \max_{a'} u(c^{jk}(z,h,a,a')) + \psi_t^{jk} + \beta \Theta_{t+1}^{jk}(z,h,a'), \tag{3.2}$$

where the additional state variables are the match quality shocks of both household members ($z = (z_1, z_2)$), their human capital levels ($h = (h_1, h_2)$), and joint asset holdings a. Households value consumption c according to the utility function $u(c)$. Consumption is pooled within the household. Additionally, instantaneous utility is affected by ψ which is allowed to depend on the labor market state and age. It captures disutility from search and the utility of staying at home. Households discount their continuation value Θ, which is described in detail below, with discount factor β.

Households choose assets for the next period subject to their budget constraint

$$c^{jk}(z,h,a,a') = \underbrace{\mathbb{I}_{j=E} w(z_1,h_1) + \mathbb{I}_{k=E} w(z_2,h_2)}_{\text{labor income}} + \underbrace{\mathbb{I}_{j=U} \bar{b} + \mathbb{I}_{k=U} \bar{b}}_{\text{unemployment benefits}} - \underbrace{(a' - (1+r)a)}_{\text{net savings}}. \tag{3.3}$$

Depending on their employment status households receive wage and benefit income. In addition to this, a household can use its assets and interest income to finance consumption and new purchases of the risk-free bond.

To write the continuation utility for one labor market state explicitly, we consider a household with two employed members today. Since both members are employed, the relevant state variables are two match quality shocks and two human capital levels. In addition, the continuation utility depends on the asset choice.

We express the continuation value in two steps. First, we take expectations over separation shocks and and the resulting choice sets for future labor market states:

$$\begin{aligned}
&\Theta_{t+1}^{EE}(z_1,z_2,h_1,h_2,a') = \\
&\quad (1-\delta(h_1))(1-\delta(h_2))\,\widetilde{V}_{t+1}(z_1,z_2,h_1,h_2,a',\mathcal{J}_{XX}^{EE}) \\
&\quad +\delta(h_1)(1-\delta(h_2))\,\widetilde{V}_{t+1}(z_1,z_2,h_1,h_2,a',\mathcal{J}_{UX}^{XE}) \\
&\quad +(1-\delta(h_1))\delta(h_2)\,\widetilde{V}_{t+1}(z_1,z_2,h_1,h_2,a',\mathcal{J}_{XU}^{EX}) \\
&\quad +\delta(h_1)\delta(h_2)\,\widetilde{V}_{t+1}(z_1,z_2,h_1,h_2,a',\mathcal{J}_{UU}^{XX}).
\end{aligned} \tag{3.4}$$

If neither member is exogenously separated (first line), both household members have the opportunity to work, but neither of them is eligible for benefits if he or she chooses to voluntarily quit. Hence, the feasible set of labor market states is denoted by $\mathcal{J}_{XX}^{EE}$. Lines 2 and 3 deal with the cases in which one member is exogenously separated whereas the last line considers the case in which both members receive the separation shock. In these instances, the exogenously separated member is eligible for benefits but cannot be employed in the next period.

In a second step, we consider transitions for match quality z and human capital h as well as the household's discrete choice over feasible future labor market states:

$$\begin{aligned}
&\widetilde{V}_{t+1}(z_1,z_2,h_1,h_2,a',\mathcal{J}_{QR}^{OP}) = \\
&\quad \phi^{up}(h_1)\phi^{up}(h_2)\,\mathbb{E}_{z_1'|z_1}\mathbb{E}_{z_2'|z_2}\,\mathbb{E}_{\epsilon}\max_{\widehat{jk}\in\mathcal{J}_{QR}^{OP}}\left\{V_{t+1}^{\widehat{jk}}(z_1',z_2',h_1+1,h_2+1,a')+\sigma\epsilon^{\widehat{jk}}\right\} \\
&\quad +\phi^{up}(h_1)(1-\phi^{up}(h_2))\,\mathbb{E}_{z_1'|z_1}\mathbb{E}_{z_2'|z_2}\,\mathbb{E}_{\epsilon}\max_{\widehat{jk}\in\mathcal{J}_{QR}^{OP}}\left\{V_{t+1}^{\widehat{jk}}(z_1',z_2',h_1+1,h_2,a')+\sigma\epsilon^{\widehat{jk}}\right\} \\
&\quad +(1-\phi^{up}(h_1))\phi^{up}(h_2)\,\mathbb{E}_{z_1'|z_1}\mathbb{E}_{z_2'|z_2}\,\mathbb{E}_{\epsilon}\max_{\widehat{jk}\in\mathcal{J}_{QR}^{OP}}\left\{V_{t+1}^{\widehat{jk}}(z_1',z_2',h_1,h_2+1,a')+\sigma\epsilon^{\widehat{jk}}\right\} \\
&\quad +(1-\phi^{up}(h_1))(1-\phi^{up}(h_2))\,\mathbb{E}_{z_1'|z_1}\mathbb{E}_{z_2'|z_2}\,\mathbb{E}_{\epsilon}\max_{\widehat{jk}\in\mathcal{J}_{QR}^{OP}}\left\{V_{t+1}^{\widehat{jk}}(z_1',z_2',h_1,h_2,a')+\sigma\epsilon^{\widehat{jk}}\right\}
\end{aligned} \tag{3.5}$$

For employed individuals human capital can either remain constant or increase. Each line of equation 3.5 corresponds to one of the resulting four combinations of possible human capital transitions. Moreover, in each case, expectations are also taken with respect to match quality shocks.

The possible choices of future labor market states can be read off Table 3.4. $\epsilon \in \mathbb{R}^{|\mathcal{J}_{QR}^{OP}|}$ is a vector of iid, Type-I extreme value (Gumbel) shocks with mean zero. We introduce these taste shocks for computational purposes, as they smooth out kinks and discontinuities in the policy functions that arise from the discrete choices over labor market states. We choose the variance of these taste shocks to be small enough such that they do not affect the solution to the problem in an economically meaningful way.

While we outline here the continuation value for a household with two members currently employed, the problem for all other current joint labor market states evolves in a very similar manner: In equation 3.4, instead of separation shocks expectations are formed over job offer arrivals and potential losses of benefit eligibility for non-employed members. Equation 3.5 remains mostly unaffected except for initial draws of z out of non-employment, which stem from an initial distribution and are independent of past realizations of z.

3.3.3 Vacancy Posting and Endogenous Arrival Rates

To determine the job arrival rates of households endogenously we consider the optimal vacancy posting problem of single-job firms. We assume free entry of firms and a cost κ of posting a vacancy. A vacancy lasts for one period and if not filled can be renewed by paying κ again.

A match with quality z between a firm and a worker with human capital h produces per period output $y(z,h)$, of which the worker receives a constant share χ as a wage $w(z,h) = \chi y(z,h)$, yielding firms' per period profit of such match as $\pi(z,h) = (1-\chi)y(z,h)$.

The expected future value to a firm of a match with a worker i from a household with current state $x_i = (t, z_i, z_{-i}, h_i, h_{-i}, a, jk)$ and asset choice for next period a' , given that the household can choose the joint future labor market state from set $\mathcal{J}_{QR}^{OP}$, is defined as

$$\begin{aligned} &EJ_{t+1}^{jk}(z_i, z_{-i}, h_i, h_{-i}, a', \mathcal{J}_{QR}^{OP}) = \\ &\qquad \mathbb{E}_{h_i'|h_i} \mathbb{E}_{h_{-i}'|h_{-i}} \mathbb{E}_{z_i'|z_i} \mathbb{E}_{z_{-i}'|z_{-i}} \mathbb{E}_{\hat{j}k \in \mathcal{J}_{QR}^{OP}} \mathbb{I}_{\hat{j}=E|x'} J_{t+1}^{\hat{j}k}(z_i', z_i', h_{-i}', h_{-i}', a') \end{aligned} \tag{3.6}$$

where $\mathbb{E}_{\hat{j}k \in \mathcal{J}_{QR}^{OP}} \mathbb{I}_{\hat{j}=E|x'}$ is firms' expectation of the household's joint labor market choice and an indicator of whether for each joint state member i stays with the firm, i.e. firms' expectation over endogenous acceptances and quits. The contemporaneous value to the firm is then given by

$$J_t^{jk}(z_i,z_{-i},h_i,h_{-i},a) = \pi(z_i,h_i) + \frac{1}{1+r}(1-\delta(h_i))\mathbb{E}_{P,R} E J_{t+1}^{jk}(z_i,z_{-i},h_i,h_{-i},a',\mathcal{J}_{XR}^{EP}), \tag{3.7}$$

where $\mathbb{E}_{P,R}$ is a firm's expectation over job loss, job finding, and eligibility transitions of the spouse and $a' = a(t,z_1,z_2,h_1,h_2,a,jk)$ is the household's asset choice.

We discuss the determination of endogenous arrival rates using the example of a household with both members unemployed but not eligible for benefits, i.e. a household with initial labor market state SS. Define member i's arrival rate as

$$\lambda_t(h_i,h_{-i},a,jk) = \lambda_S p(\theta_t(h_i,h_{-i},a,jk)) \tag{3.8}$$

with arrival rate $p(\theta) = m(1,\theta)$ and corresponding vacancy filling rate $q(\theta) = m(\frac{1}{\theta},1)$, where $m(U,V)$ is the standard Cobb-Douglas matching function, with market tightness θ denoting the ratio of vacancies over searchers in any given submarket. Hence $p(\theta) = \theta^{1-\alpha}$, $q(\theta) = \theta^{-\alpha}$, and $p(\theta) = \theta q(\theta)$. λ_S is an exogenous shifter that only depends on the previous labor market state and reflects the consequences of differences in search effort between unemployed (U or S) and out of the labor force (N). This distinction is necessary because – conditional on the remaining states of the household – firms will not differentiate whether they hire a worker out of unemployment or from out of the labor force.

Free entry imposes that the expected value of a vacancy (probability of filling times the value if filled) has to equal the cost of posting κ. This condition determines relevant market tightness $\theta_t(h_i,h_{-i},a,jk)$. The free entry condition needs to satisfy

$$\kappa = q(\theta_t(h_i,h_{-i},a,jk))\mathbb{E}_P E J_{t+1}^{jk}(z_i,z_{-i},h_i,h_{-i},a',\mathcal{J}_{XX}^{EP}). \tag{3.9}$$

Here $\mathbb{E}_P$ captures expectations over the spouse's job finding and is an equation in the spouse's $\theta_t(h_{-i},h_i,a,jk)$ as the spouse is also currently not employed. Hence, in all cases with currently two non-employed household members we have to solve a system of two non-linear equations in two unknowns.

With slight abuse of notation the two equations solving for two θs can be written as

$$\kappa = q(\theta_i)[\lambda(\theta_{-i})\underbrace{EJ^{SS}_{t+1}(h_i,h_{-i},a',\mathcal{J}^{EE}_{XX})}_{EJ^{EE}_i} + (1-\lambda(\theta_{-i}))\underbrace{EJ^{SS}_{t+1}(h_i,h_{-i},a',\mathcal{J}^{EX}_{XX})}_{EJ^{EX}_i}], \tag{3.10}$$

$$\kappa = q(\theta_{-i})[\lambda(\theta_i)\underbrace{EJ^{SS}_{t+1}(h_{-i},h_i,a',\mathcal{J}^{EE}_{XX})}_{EJ^{EE}_{-i}} + (1-\lambda(\theta_i))\underbrace{EJ^{SS}_{t+1}(h_{-i},h_i,a',\mathcal{J}^{EX}_{XX})}_{EJ^{EX}_{-i}}]. \tag{3.11}$$

This yields

$$\theta_{-i} = \left[\frac{\kappa}{\lambda(\theta_i)EJ^{EE}_{-i} + (1-\lambda(\theta_i))EJ^{EX}_{-i}}\right]^{-\frac{1}{\alpha}} \tag{3.12}$$

and hence

$$\begin{aligned}\kappa = q(\theta_i)&\left[\lambda_S\left[\frac{\kappa}{\lambda(\theta_i)EJ^{EE}_{-i} + (1-\lambda(\theta_i))EJ^{EX}_{-i}}\right]^{\frac{\alpha-1}{\alpha}} EJ^{EE}_i\right.\\ &+\left.\left(1-\lambda_S\left[\frac{\kappa}{\lambda(\theta_i)EJ^{EE}_{-i} + (1-\lambda(\theta_i))EJ^{EX}_{-i}}\right]^{\frac{\alpha-1}{\alpha}}\right) EJ^{EX}_i\right],\end{aligned} \tag{3.13}$$

which is a non linear equation in one unknown and can be solved numerically.

The endogenous arrival rates can be derived in a similar fashion for other cases of original labor market states. The exogenous component of λ needs to be adjusted to reflect whether an agent is unemployed or out of the labor force. Solving for endogenous arrival rates gets substantially easier if one spouse has been previously employed since in this case we only have one θ and hence we only need to solve one equation with one unknown.

Given this setup, job finding probabilities of an individual depend on all the state variables, including assets, age, and own human capital, but also the spouse's human capital, employment status, and potentially match quality. With regard to age, our setup is hence able to capture that it may be harder for older workers to find a new job. In the model, firms are less willing to hire older workers because they have to retire at a certain age, leaving less time to recover the vacancy posting cost. In our calibration, this effect is strong close to retirement but relatively weak at young ages because in these cases it is quite likely that the match is dissolved before retirement in any case.

It is also intuitive that arrival rates depend on an individual's human capital. It is potentially less appealing that we also condition on the spouse's state variables. It is necessary, however, because it influences the probabilities of an individual accepting a certain job and quitting later on. Having different submarkets and free entry in each active submarket simplifies computation drastically, as we do not need to know the distribution of individuals across states to solve for arrival rates.

This setup for determining age-dependent arrival rates in the labor market generally implies arrival rates decreasing in age, decreasing in assets because richer individuals are more likely to quit, increasing in human capital because the value of the match is higher and individuals are less likely to quit, increasing in match quality for the same reasons, and decreasing in a spouse's employment, human capital, and match quality because having a spouse earning high wages increases the quit probability and lowers the value of a match to the firm.

3.3.4 Numerical Implementation

In our setup, agents do not face risk during retirement. This assumption renders the household problem during retirement very simple. We solve the retirement problem using the endogenous grid method of Carroll (2006) to obtain a terminal condition for the household problem during working life.

The household problem during working life is high-dimensional because of the many combinations of labor market states and the fact that we have to keep track of match quality shocks and human capital for both members. Furthermore, given our focus on labor market transitions, the model has a monthly frequency. For computational efficiency, we therefore solve the household problem following Iskhakov et al. (2017), who extend the endogenous grid-point method of Carroll (2006) to problems with discrete and continuous choices. Thus, their approach is well suited for our problem with a discrete choice over labor market states and a continuous asset choice.

The algorithm proceeds as follows: Within each period, given future value functions of both the household and firm, we begin by determining households' choices over future labor market states for each potential choice set. With this, we are able to solve firms' vacancy posting problem and determine endogenous arrival rates. Endogenous arrival rates given, we can solve households' consumption-savings problem as described above. In a final step, we update households' and firms' value functions making use of households' policy functions and again the endogenous arrival rates.

3.4 Calibration

We solve the model at a monthly frequency. This assumption is in line with the frequency at which we observe labor market transitions in the data and necessary because the U.S. labor market exhibits high rates of turnover. We assume that the period of working life is 40 years, corresponding to 480 months. The retirement period is another 120 months, i.e. 10 years.

3.4.1 Functional Form Assumptions

Households value consumption with a standard CRRA utility function

$$u(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}, \tag{3.14}$$

where γ is the coefficient of relative risk aversion. The second part of instantaneous utility that has to be parameterized is the parameter ψ_t^{jk} which differs across joint labor market states, reflecting disutility of work and search. Furthermore, we allow it to vary by age.[8]

Output is assumed to be the product of human capital and the match quality shock:

$$y(h,z) = hz. \tag{3.15}$$

Human capital is defined on an equidistant grid. The probabilities of moving to a higher (lower) human capital level when employed (non-employed) are given by the following processes:

$$\phi^{up}(i) = \bar{\phi}^{up} i^{\underline{\phi}^{up}} \tag{3.16}$$

$$\phi^{down}(i) = \bar{\phi}^{down} i^{\underline{\phi}^{down}}, \tag{3.17}$$

where i indicates the grid point rather than the level of human capital. This process is able to capture falling or rising probabilities of moving up or down the human capital ladder. The match quality shock while employed is assumed to follow an autoregressive process of order 1 in logs. We discretize the process using the method of Tauchen (1986).

Finally, we have to make an assumption on the arrival rates of job offers and separation rates in the labor market. We restrict $\lambda_S, \lambda_U, \lambda_N$ to be constant across age.[9] We allow the separation rate to vary with human capital according to a similar process as the probabilities of moving up or down the human

[8] In the current calibration, the disutility of search parameter is mostly constant across age. In fact, we make an exception only for one labor market state, as discussed below.

[9] Even though the exogenous component of arrival rates is constant in age, the solution to firms' vacancy posting problem endogenously yields arrival rates falling in age t conditional on households' remaining states.

capital ladder:

$$\delta(i) = \bar{\delta} i^{\underline{\delta}}. \tag{3.18}$$

3.4.2 Parameters and Moments

To compare the model to the data, we simulate the full life cycle of 40,000 households and compute model-implied moments of this simulation. We initialize the distribution of households across labor market states such that it is consistent with the data. We assume that all agents start with one of the lowest asset levels. For employed individuals, we draw the match quality shock from the stationary distribution of the match quality process. For human capital, even though this is mostly supposed to capture work experience in our model, we assume some heterogeneity in the initial distribution to obtain sufficient dispersion in incomes. Human capital levels are, however, concentrated on the lower rungs of the human capital ladder.

While in the model all parameters jointly determine all moments, we now discuss which parameters are most closely related to which moments. Table 3.5 summarizes the parameter values. We start by setting a number of parameters without solving the model. We exogenously fix the coefficient of relative risk aversion to two, a standard value in the literature. We set the monthly net interest rate to 0.17%, corresponding to an annual interest rate of roughly 2%. We assume a probability of losing unemployment benefits of $\phi^{US} = 1/6$, consistent with an average duration of benefit receipt of six months. Finally, we set the elasticity of the matching function α to 0.5, as in Petrongolo and Pissarides (2001), and the share of match output going to the worker χ to 0.7.

We target key moments of the U.S. labor market that are related to a large number of parameters. First, we target individual transition rates between labor market states. These are closely related to the parameters $\lambda_N, \lambda_S, \lambda_U$, the exogenous upper bounds on arrival rates depending on labor market states. We impose the restriction $\lambda_S = \lambda_U$, as these two states only differ in whether an individual receives unemployment benefits or not. Individual transition rates are closely related to the vacancy posting cost κ. The EU rate in particular pins down parameters of the job loss process. The model captures well the magnitude of the transitions between employment and unemployment. In contrast, it undershoots the magnitude of transitions between non-employment and employment/unemployment, as we will discuss in more detail in the next section when looking at the added worker effect in the model.

Another important set of targeted labor market moments is the distribution of households over joint labor market states for four ten-year age groups. Because the arrival rates are endogenously determined from the firm problem we treat the preference parameters ψ that govern the disutility of work and search as free parameters to match joint labor market states by age. We keep all these parameters constant by age, except for $\psi^{EN} = \psi^{NE}$, which we assume to be decreasing with age. Specifically, we assume $\psi^{EN} = \psi^{NE}$ to start at a level of 1.30 at age 25 and to decay logistically to a level of 90 with a half-life of 100 months. Imposing this age-dependency is necessary in order to avoid that too many young households have both members employed. Economically, we justify a higher utility of having one member at home for young households because this is the age group who are most likely to have young children. As we do not model children explicitly, introducing age-dependency in $\psi^{EN} = \psi^{NE}$ is a parsimonious way of capturing this motive and helps us to match a high enough share of young households with one member employed and one member out of the labor force.

In addition to these labor market moments, we target life cycle profiles of income and assets. The pension level p and the discount factor β are mostly determined by the shape of the life cycle asset profile. Specifically, we target mean asset holdings for four age groups. An important question is which assets to consider in the data when constructing the moments to be matched. For insurance reasons, the relevant concept is liquid assets. In particular, because a model period is one month, it would be desirable to consider only assets that can be liquidated at a monthly frequency. However, given the life cycle dimension of our setup, retirement is an important driver of savings. Imposing too strict requirements on asset liquidity would exclude much of households' retirement savings. Therefore, considering the trade-off between asset liquidity and retirement savings, we choose to target financial assets including retirement accounts net of debt. In addition, we include vehicle equity because it can be accessed very quickly. In contrast, we exclude houses and mortgages because tapping into home equity is difficult for unemployed and might take longer, so it is not as useful for insurance purposes on a monthly frequency. Business equity is excluded for the same reason. We construct asset-related data moments from the Panel Study of Income Dynamics (PSID).

The parameters of the human capital process are chosen to match the income profile over the life cycle. In the data, these moments are constructed from the PSID. The probability of moving up the human capital ladder is decreasing in the human capital level which is a way of achieving a concave income profile: When young, an agent moves up the human capital ladder quickly such that the wage increase is steeper. After a few steps on the human capital ladder, the likelihood of a further increase in human capital decreases quite significantly such that the income profile becomes flatter. The probability of losing human capital, by contrast, is constant across human capital levels. Human capital decay of non-employed allows us to capture the empirical observation that newly employed individuals have lower wages than long-time employed and that job losses lead to persistent wage drops (Davis and von Wachter, 2011; Jarosch, 2015; Kospentaris, 2021).

The parameters of the match quality shock process are chosen to match the variance in income levels by age group. Additionally, we have to pin down the distribution from which newly employed draw their match quality, which we set to the stationary distribution of the discretized Markov chain.

The only remaining parameters to be set are the level of the unemployment benefit and the variance of the taste shock. We assume the unemployment benefit to be constant and set its level to be roughly 50% of median income. For the taste shock, we set $\sigma_{\varepsilon} = 0.1$. Using 0.05 instead does not meaningfully impact or results.

3.4.3 Fit of Targeted Moments

In this section, we present the model fit for key targeted moments. First, Figure 3.6 shows the share of households in joint labor market states by age group in the model and in the data. To compare the model to the data, we pool all agents who are unemployed with and without benefits into one group, labeled U. In all age groups, the most common joint labor market state is that both members are employed. This share is, however, strongly decreasing in age, with around 65% of households being in that group among the two young groups and just 45% in the oldest age group. By contrast, the share of households where at least one member is out of the labor force is increasing over the life cycle. Among the youngest there are very few households with both members out of the labor force. Among the oldest, almost 20% of all couples are jointly non-participating. In addition, the share of households with one member employed and one member out of the labor force is slightly increasing in age. Overall, the model matches very well the distribution of households over joint labor market states. It also captures that the share of two earner households is decreasing in age and that the share of households with at least one member out of the labor force is increasing in age, though it somewhat understates the magnitude of these changes over the life cycle.

Moreover, the model is able to replicate average asset holdings over the life cycle, as shown in Table 3.6. Averaging over all age groups, we match the average asset level of the population well. However, the model slightly underpredicts the mean asset holdings of the medium age groups. However, it captures that average asset holdings are strongly increasing in age.

Figure 3.6 Joint Labor Market States of Couples (Model vs. Data)

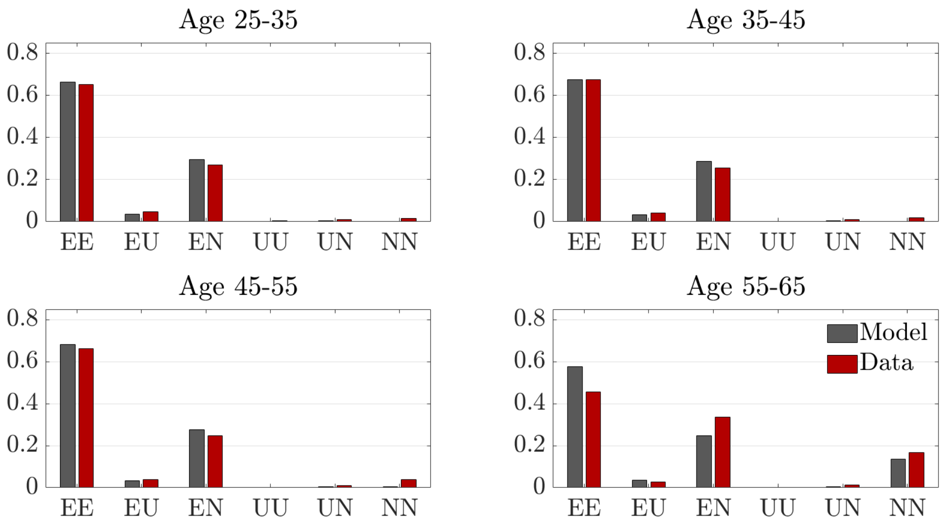

Notes: Figure 3.6 shows the joint labor market states of couples in the model and in the data. For the model, U includes both unemployed receiving benefits and searchers who do not receive benefits. The data is from the CPS.

Finally, we consider the model fit for mean income levels and the dispersion in income across age groups. Table 3.7 shows the comparison between data and model. Again, when averaging over all age groups, the model is close to the income level in the data but as of now undershoots the dispersion. Moreover, the model is able to replicate the increase in mean income for the age groups 25-35, 35-45, and 45-55. It fails, however, in generating a fall in income for the oldest group. This mismatch for the oldest age group arises from a strong selection effect in the model with respect to who stays in the labor force. Many agents with relatively low human capital and/or match quality prefer to drop out of the labor force, which drives up the average income among the employed. In contrast, the model replicates that income dispersion within age group is higher among the old than among the young.

Table 3.5 Parameter Values

Parameter	**Interpretation**	**Value**
Demographics		
T	Length of life in months	600
T_W	Length of working life in months	480
Preferences		
β	Discount factor	0.9955
γ	Risk aversion	2.0000
$\psi^{EE}, \psi^{EU}, \psi^{UE}, \psi^{ES}, \psi^{SE}$	Disutility of work/search	0.0000
$\psi^{UU}, \psi^{SS}, \psi^{SU}, \psi^{US}$	Disutility of work/search	0.5000
$\psi^{UN}, \psi^{NU}, \psi^{SN}, \psi^{NS}$	Disutility of work/search	1.2000
ψ^{NN}	Disutility of work/search	2.6000
ψ^{EN}, ψ^{NE}	Disutility of work/search	$1.3 + \frac{0.9-1.3}{1+e^{-0.05(t-100)}}$
Financial Assets		
r	Interest rate	0.0017
Labor Market		
$\bar{\delta}$	Level parameter separation rate	0.0200
$\underline{\delta}$	Curvature parameter separation rate	-0.5000
λ_U, λ_S	Probability of job offer for unemployed	0.4500
λ_N	Probability of job offer out of labor force	0.3000
Human Capital		
$\underline{h}$	Lower bound h	0.2000
$\bar{h}$	Upper bound h	0.8000
$\bar{\phi}^{up}$	Level parameter prob. h rise	0.0500
$\underline{\phi}^{up}$	Curvature parameter prob. h rise	-1.2000
$\bar{\phi}^{down}$	Level parameter prob. h fall	0.3316
$\underline{\phi}^{down}$	Curvature parameter prob. h fall	0.0000
Match Quality Shocks		
ρ_z	Persistence	0.9000
σ_z	Standard deviation	0.1000
Firms		
χ	Labor share of output	0.7000
κ	Cost of vacancy posting	8.0000
α	Matching elasticity	0.5000
Government		
b	Unemployment benefit	0.2500
ϕ^{US}	Probability of losing benefits	0.1667
p	Pension	0.2000
Gumbel shock		
σ_ε	Standard deviation of taste shock	0.1000

Notes: Table 3.5 summarizes the parameter values.

Table 3.6 Asset Levels

	Model	Data
All	10.4	11.8
Age 25-35	2.8	3.0
Age 36-45	4.9	7.0
Age 46-55	10.6	14.6
Age 55-65	23.3	24.1

Notes: Table 3.6 compares mean asset holdings by age group in the model and in the data. The data is from the PSID. In the data, assets include financial assets net of debt and vehicle equity. 1 unit corresponds to $10,000.

Table 3.7 Income Levels and Dispersion

	Level		Standard deviation	
	Model	Data	Model	Data
All	0.3596	0.3424	0.1363	0.2374
Age 25-35	0.3296	0.3020	0.1172	0.2009
Age 36-45	0.3538	0.3572	0.1341	0.2456
Age 46-55	0.3752	0.3629	0.1429	0.2486
Age 56-65	0.3826	0.3400	0.1511	0.2466

Notes: Table 3.7 compares mean and standard deviation of labor income by age group in the model and in the data. The data is from the PSID. 1 unit corresponds to $10,000.

3.5 Results

In this section we first present the model implications for untargeted moments. Second, we show that our model can replicate the decreasing magnitude of the added worker effect over the life cycle. Third, we use the model to construct counterfactuals and analyze which channels are responsible for the age-dependency in the added worker effect.

3.5.1 Untargeted Moments

We begin this section by presenting untargeted life cycle profiles of individual labor market transitions in Figure 3.7. Again, in the model U comprises both the group of unemployed who receive benefits and those who exert costly search effort without receiving benefits.

First, consider transitions from employment over the life cycle (Figure 3.7a to 3.7c). The model captures that the likelihood of remaining in employment falls quite rapidly towards the end of working life, though the monthly transition probability out of employment never falls below 95%. The counterpart to this in model and data is a corresponding increase in the likelihood of moving from employment to out of the labor force. As agents get closer to the retirement age, it is not worthwhile for them to stay employed when they receive a bad match quality shock or have low human capital. By contrast, young agents continue to work even in these cases. Several model mechanisms account for this. First, young agents have a longer time horizon until retirement, so that they need labor income to cover consumption needs during working life. In contrast, old agents hold much higher levels of assets which they can use to finance consumption. Second, human capital is only accumulated while employed. Thus, higher human capital is more valuable for the young as they can benefit from it for a longer time period. The model performs very well in matching the slightly decreasing path of E to U transitions over the life cycle.

Next, consider the transitions out of unemployment (Figure 3.7d to 3.7f). The model replicates that across the entire life cycle the most likely transition is to remain unemployed. It also matches well that the probability of transitioning to employment declines with age, whereas the probability of giving up on searching and leaving the labor force increases with age. Finally, the model generates a fall in transitions from out of the labor force into employment (Figure 3.7g) but understates the likelihood to transition into unemployment (Figure 3.7h) over the life cycle, while it matches well the high persistence of non-participation (Figure 3.7i).

Figure 3.7 Labor Market Transitions over the Life Cycle

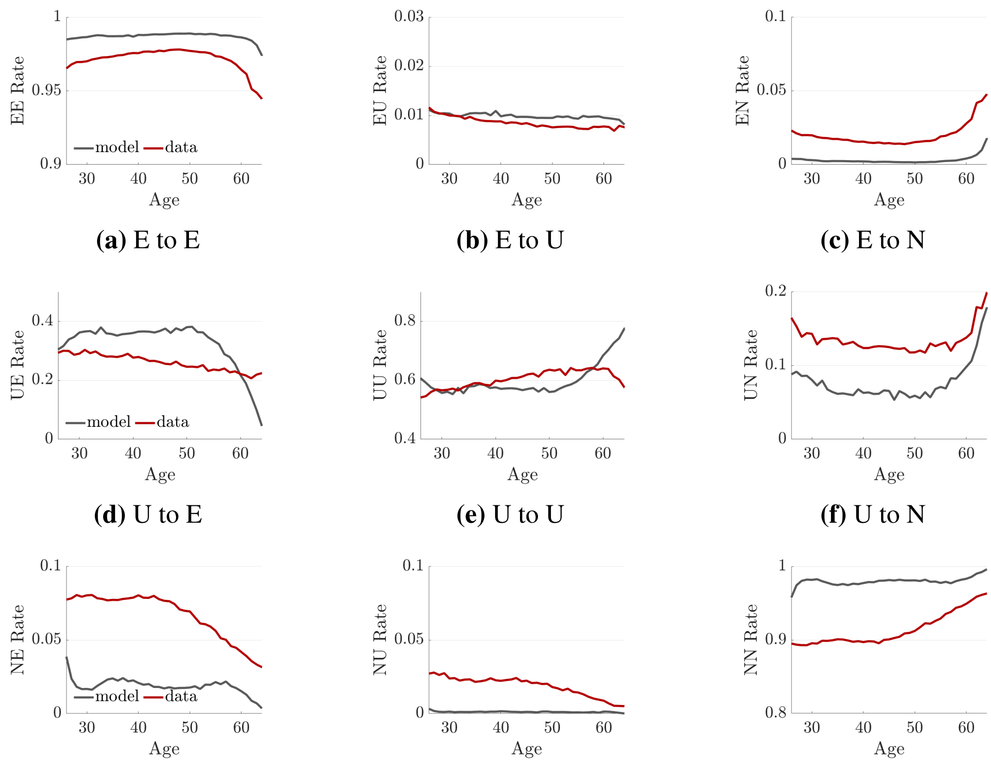

(a) E to E **(b)** E to U **(c)** E to N

(d) U to E **(e)** U to U **(f)** U to N

(g) N to E **(h)** N to U **(i)** N to N

Notes: Figure 3.7 shows individual labor market transitions in the data and in the model. For the model, U includes both unemployed receiving benefits and searchers who do not receive benefits. The data is from the CPS.

Again, it is apparent from these figures that the model generates too few transitions between out of the labor force and employment/unemployment. This is most likely due to the fact that we leave many important life events such as child birth, marital transitions, and health shocks unmodeled. We will show next, however, that the model captures well the impact of one key life event, job loss of the primary earner, on the labor force participation of out of the labor force spouses, the added worker effect.

3.5.2 The Added Worker Effect over the Life Cycle in the Model

We now evaluate whether the model can replicate our main empirical finding: the age dependency in the added worker effect. To compare model to data, we replicate Table 3.3 from Section 3.2 with simulated model data in Table 3.8. For ease of comparison, we also report empirical transition probabilities.

For the young, the model is capable of producing a strong increase in the probability of moving from out of the labor force directly into employment and into unemployment upon job loss of the primary earner. The model generally underestimates the probability of spousal transitions directly into employment

Table 3.8 Joint Labor Market Transitions by Age (Model vs. Data)

	Primary earner transition	
	EE	EU/ES
Young (25-35):		
Cond. prob. of spousal NE transition	2.26%	3.12%
	6.66%	9.30%
Cond. prob. of spousal NS transition	0.40%	5.28%
	2.00%	6.89%
Cond. prob. of spousal NN transition	97.34%	91.60%
	91.34%	83.81%
Old (55-65):		
Cond. prob. of spousal NE transition	1.95%	2.24%
	4.29%	3.73%
Cond. prob. of spousal NS transition	0.11%	1.16%
	0.90%	2.75%
Cond. prob. of spousal NN transition	97.95%	96.60%
	94.81%	93.52%

Notes: This table compares joint labor market transitions by age in the model and in the data.

independently of the primary earner's transition. However, it captures very well the difference in probabilities depending on the primary earner transition, which is the added worker effect.

In the model, as in the data, there is a much smaller added worker effect for the old. The model reproduces that there is no substantially increased likelihood of transitioning from out of the labor force directly into employment when the primary earner loses a job for the old. Furthermore, the increased probability of searching for a job by exerting costly effort is much lower than for the young, in line with the data.

Hence, the model performs well in generating the instantaneous added worker effect over the life cycle. To analyze anticipation effects and lagged responses, Figure 3.8 replicates Equation (3.1) on model simulated output, separately by age. In line with the data, the model produces larger contemporaneous and lagged effects for the young than for the old. The lead effects are, however, of similar size across both age groups.

The model mechanisms that produce lagged responses are threefold. First, after becoming unemployed the primary earner may lose human capital which decreases potential human capital differences across spouses. Consequently, it may be optimal that both spouses search or to re-optimize on the actively searching household member. Second, unemployment benefits can expire, making employment a more desirable state. Third, households without any employed member may run down their assets to finance

Figure 3.8 Dynamic Response: AWE by Age in the Model

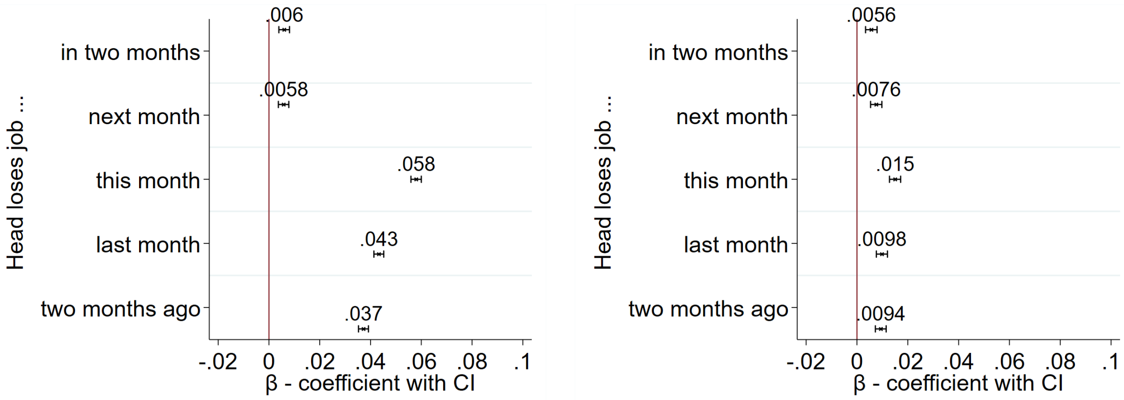

(a) Young: 25-35 years **(b)** Old: 55-65 years

Notes: Figure 3.8 shows the change in the probability that a non-participating spouse enters the labor force (either as unemployed or as employed) this month if household head loses/lost the job in two months, next month, this month, last month, two months ago, respectively, relative to the baseline in which the household head remains employed. Figure 3.8a shows the model results for young households; Figure 3.8b shows the model results for old households. The regression producing the coefficients is Equation (3.1).

consumption, which increases the need to search for a new job to re-accumulate assets for precautionary reasons and for retirement.

While the model produces some anticipation effect in the two months prior to a primary earner's job loss, these lead effects are smaller than in the data. Job loss is predictable because the exogenous separation probability depends on human capital. Spouses of low human capital employed individuals may enter the labor force because a future separation is relatively likely, whereas spouses of high human capital individuals choose not to do so because the chance of an exogenous separation is low. By the law of large numbers, these separations do in fact realize at higher rates for low human capital primary earners, producing the effect that spouses are more likely to enter the labor force in anticipation of a job loss. In addition, persistence in match quality might induce non-participating spouses to enter the labor force upon a decline in match quality for the employed spouse, preparing a potential future quit if match quality remains low.

3.5.3 Counterfactuals

Finally, we use the model to construct counterfactuals and analyze which channels are important in driving the age-dependency in the added worker effect. For that purpose, we start with the added worker effect of the young and then change individual model elements towards the counterparts of old households. Table 3.9 reports the results for three such counterfactuals together with the baseline results for young households.

Table 3.9 Joint Labor Market Transitions Counterfactuals

	Primary earner transition	
	EE	EU/ES
Young (25-35):		
Cond. prob. of spousal NE transition	2.26%	3.12%
Cond. prob. of spousal NS transition	0.40%	5.28%
Cond. prob. of spousal NN transition	97.34%	91.60%
Counterfactual meeting probabilities		
Cond. prob. of spousal NE transition	2.14%	2.93%
Cond. prob. of spousal NS transition	0.41%	5.36%
Cond. prob. of spousal NN transition	97.46%	91.71%
Counterfactual human capital		
Cond. prob. of spousal NE transition	1.70%	3.02%
Cond. prob. of spousal NS transition	0.24%	3.09%
Cond. prob. of spousal NN transition	98.06%	93.89%
Counterfactual assets		
Cond. prob. of spousal NE transition	0.11%	0.33%
Cond. prob. of spousal NS transition	0.11%	0.43%
Cond. prob. of spousal NN transition	99.78%	99.23%

Notes: This table shows the counterfactual joint labor market transition probabilities.

The first counterfactual adjusts job arrival rates for young households. More specifically, we first compute the average job arrival rate for old and for young households in the model, restricting the sample to households with one member employed and one member out of the labor force. Afterwards, we adjust the individual arrival rates of each young household in our simulation by the difference between these previously computed means. This approach moves the average arrival rate of young households to that of their old counterparts, but preserves the relative distribution of arrival rates among the young. The second block of Table 3.9 shows that adjusting arrival rates has a limited impact on the added worker effect. This result arises because the average arrival rates for young and old are very similar: As most non-participating spouses are unlikely to accept a job offer, firms are only offering low arrival rates in order to satisfy their free entry condition. Nevertheless, the average arrival rate is slightly lower for older households resulting in fewer employment transitions both in the EE and in the EU case.[10]

[10] This result may be partially due to the timing assumptions in the model. At the moment firms post vacancies in all the submarkets before separation shocks occur. Hence, out of the labor force spouses do not consider

In the second counterfactual, we adjust the human capital level of young households. Similar to above, we compute the difference in mean human capital levels across age groups separately for employed and non-participating spouses and adjust the human capital level of each young household by the difference. In our simulation, the employed spouse among older households has a higher human capital due to on average longer cumulative employment spells. In contrast, human capital levels for non-participating spouses are very similar across age groups. This is partially driven by selection (low human capital individuals are more likely to be non-participating when they have an employed spouse) and partially by fast depreciation of human capital during non-employment in order to match empirical wage losses from non-employment spells. Thus, the results of the second counterfactual can be attributed to a higher human capital level of the employed spouse during old age.

The third block of Table 3.9 shows that the increase in human capital of the employed spouse reduces transition probabilities into participation for both the EE and the EU case, but also dampens the added worker effect. When the human capital of a separated spouse is higher, this spouse is more likely to find a new job (arrival rates are increasing in human capital) and the difference in human capital levels across spouses is potentially larger, making a switch in the prime earner position less likely.

In a third counterfactual, we adjust the asset levels of young households in the same manner as arrival rates and human capital. Since old households have on average substantially higher asset levels we make all young households richer. The fourth block of Table 3.9 shows that this eliminates the incentive for a non-participating spouse to transition into participation. Hence, the added worker effect vanishes. Young households with asset holdings of the old are relatively rich for their age, reducing the incentive to work also in the baseline EE case, and are well insured against any labor market shock such that they do not have to rely on the added worker effect as a margin of insurance.

Taking all three counterfactuals together, we find that the substantially lower added worker effect among the old predominantly arises through higher wealth levels. Hence, older households exhibit a weaker AWE because they have better access to self-insurance through savings and are therefore less in need of other insurance margins, as opposed to a lack of opportunity to make use of the AWE.

that their partner loses the job, translating into low acceptance probabilities and in turn low vacancy posting rates. In future work, we will investigate the robustness of the finding to different timing assumptions.

3.6 Conclusion

In this paper, we provide evidence that the added worker effect is an important insurance margin against job loss of the primary earner for two-member households, but that the prevalence of this insurance channel strongly differs over the life cycle. When the primary earner transitions from employment to unemployment, an out of the labor force spouse is much more likely to enter the labor force in order to offset the income loss compared to when the primary earner remains employed. In particular, this spousal labor supply response is very strong for young households and becomes continuously weaker as households age.

To analyze the mechanisms that drive this age-dependency, we build a stochastic life cycle model of two-member households with a frictional labor market. We calibrate the model economy to match salient features of the US labor market. The model endogenously generates the added worker effect and its decreasing magnitude over the life cycle. Model counterfactuals reveal that the added worker effect is weaker for old than for young households mainly because older households are better insured through larger asset holdings, so that their need for spousal insurance is lower. In addition, human capital of employed spouses is higher for the old, making the spousal labor supply less valuable, though this channel is quantitatively smaller. Differences in arrival rates across age groups contribute little to the difference in the added worker effect due to a general reluctance of firms to offer jobs to non-participating workers.

References

Adamoupoulou, E., Hannusch, A., Kopecky, K., and Obermeier, T. (2021). Rising cohabitation and child development.

Adelino, M., Schoar, A., and Severino, F. (2021). Perception of house price risk and homeownership. Technical report.

Agnew, J., Balduzzi, P., and Sunden, A. (2003). Portfolio choice and trading in a large 401 (k) plan. *American Economic Review*, 93(1):193–215.

Alan, S. (2006). Entry costs and stock market participation over the life cycle. *Review of Economic Dynamics*, 9(4):588–611.

Almenberg, J. and Dreber, A. (2015). Gender, stock market participation and financial literacy. *Economics Letters*, 137:140–142.

Ameriks, J. and Zeldes, S. P. (2004). How do household portfolio shares vary with age. Technical report, working paper, Columbia University.

Andersen, S., Marx, J., Nielsen, K. M., and Vesterlund, L. (2020). Gender differences in negotiation: Evidence from real estate transactions. *NBER WP No 27318*.

Attanasio, O. P., Bottazzi, R., Low, H. W., Nesheim, L., and Wakefield, M. (2012). Modelling the demand for housing over the life cycle. *Review of Economic Dynamics*, 15(1):1–18.

Barber, B. M. and Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics*, 116(1):261–292.

Bardóczy, B. (2020). Spousal insurance and the amplification of business cycles. *Working Paper*.

Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–64.

Bartscher, A. (2020). It takes two to borrow: The effects of the equal credit opportunity act on homeownership and mortgage debt of married couples.

Benabou, R. (2002). Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica*, 70(2):481–517.

Bertocchi, G., Brunetti, M., and Torricelli, C. (2011). Marriage and other risky assets: a portfolio approach. *Journal of Banking & Finance*, 35(11):2902–2915.

Birinci, S. (2019). Spousal labor supply response to job displacement and implications for optimal transfers. *Working Paper*.

Blackburn, M. L. (2010). Internal migration and the earnings of married couples in the united states. *Journal of Economic Geography*, 10(1):87–111.

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *The American Economic Review*, 106(2):387–435.

Bogan, V. L. and Fernandez, J. M. (2017). How children with mental disabilities affect household investment decisions. *American Economic Review*, 107(5):536–40.

Borella, M., De Nardi, M., and Yang, F. (2018). The aggregate implications of gender and marriage. *The Journal of the Economics of Ageing*, 11:6–26.

Borella, M., De Nardi, M., and Yang, F. (2019). The lost ones: the opportunities and outcomes of non-college educated americans born in the 1960s. *NBER WP No 25661*.

Brandsaas, E. E. (2021). Household stock market participation and exit: The role of homeownership.

Bredtmann, J., Otten, S., and Rulff, C. (2018). Husband's unemployment and wife's labor supply: the added worker effect across Europe. *ILR Review*, 71(5):1201–1231.

Burdett, K. (1978). A theory of employee job search and quit rates. *The American Economic Review*, 68(1):212–220.

Burke, J. and Miller, A. R. (2018). The effects of job relocation on spousal careers: Evidence from military change of station moves. *Economic Inquiry*, 56(2):1261–1277.

Cagetti, M. (2003). Wealth accumulation over the life cycle and precautionary savings. *Journal of Business & Economic Statistics*, 21(3):339–353.

Campbell, J. Y. (2006). Household finance. *The journal of Finance*, 61(4):1553–1604.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.

Catherine, S. (2019). Countercyclical labor income risk and portfolio choices over the life-cycle. *HEC Paris Research Paper No. FIN-2016-1147*.

Catherine, S. (2020). Countercyclical labor income risk and portfolio choices over the life-cycle. *HEC Paris Research Paper No. FIN-2016-1147*.

Chang, M. (2019). A house without a ring: The role of changing marital transitions for housing decisions.

Chéron, A., Hairault, J.-O., and Langot, F. (2011). Age-dependent employment protection. *The Economic Journal*, 121(557):1477–1504.

Chéron, A., Hairault, J.-O., and Langot, F. (2013). Life-cycle equilibrium unemployment. *Journal of Labor Economics*, 31(4):843–882.

Chetty, R., Sándor, L., and Szeidl, A. (2017). The effect of housing on portfolio choice. *The Journal of Finance*, 72(3):1171–1212.

Choi, J., Zhu, J., Goodman, L., Ganesh, B., and Strochak, S. (2018). Millennial homeownership: why is it so low, and how can we increase it. *Urban Institute. https://www. urban. org/sites/default/files/publication/98729/millennial_homeownership_0. pdf*.

Choi, S. and Valladares-Esteban, A. (2020). On households and unemployment insurance. *Quantitative Economics*, 11(1):437–469.

Christelis, D., Georgarakos, D., and Haliassos, M. (2013). Differences in portfolios across countries: Economic environment versus household characteristics. *The Review of Economics and Statistics*, 95(1):220–236.

Christiansen, C., Joensen, J. S., and Rangvid, J. (2015). Understanding the effects of marriage and divorce on financial investments: the role of background risk sharing. *Economic Inquiry*, 53(1):431–447.

Cocco, J. F. (2005). Portfolio choice in the presence of housing. *The Review of Financial Studies*, 18(2):535–567.

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

Cooper, R. and Zhu, G. (2016). Household finance over the life-cycle: What does education contribute? *Review of Economic Dynamics*, 20:63–89.

Croson, R. and Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, 47(2):448–74.

Cubeddu, L. and Ríos-Rull, J.-V. (2003). Families as shocks. *Journal of the European Economic Association*, 1(2-3):671–682.

Cullen, J. B. and Gruber, J. (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics*, 18(3):546–572.

Davis, M. A., Lehnert, A., and Martin, R. F. (2008). The rent-price ratio for the aggregate stock of owner-occupied housing. *Review of Income and Wealth*, 54(2):279–284.

Davis, M. A. and Ortalo-Magné, F. (2011). Household expenditures, wages, rents. *Review of Economic Dynamics*, 14(2):248–261.

Davis, S. J. and von Wachter, T. (2011). Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, pages 1–73.

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

De Nardi, M., Fella, G., Knoef, M., Paz-Pardo, G., and Van Ooijen, R. (2021a). Family and government insurance: Wage, earnings, and income risks in the netherlands and the us. *Journal of Public Economics*, 193:104327.

De Nardi, M., Fella, G., and Paz-Pardo, G. (2019). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association*.

De Nardi, M., French, E., and Jones, J. B. (2010). Why do the elderly save? the role of medical expenses. *Journal of Political Economy*, 118(1):39–75.

De Nardi, M., French, E., Jones, J. B., and McGee, R. (2021b). Why do couples and singles save during retirement?

DiPasquale, D. and Glaeser, E. L. (1999). Incentives and social capital: Are homeowners better citizens? *Journal of urban Economics*, 45(2):354–384.

Ellieroth, K. (2019). Spousal insurance, precautionary labor supply, and the business cycle-a quantitative analysis. *Working Paper*.

Elsby, M. W. L., Hobijn, B., and Şahin, A. (2015). On the importance of the participation margin for labor market fluctuations. *Journal of Monetary Economics*, 72:64–82.

Erosa, A. and Gervais, M. (2002). Optimal taxation in life-cycle economies. *Journal of Economic Theory*, 105(2):338–369.

Fagereng, A., Gottlieb, C., and Guiso, L. (2017). Asset market participation and portfolio choice over the life-cycle. *The Journal of Finance*, 72(2):705–750.

Fagereng, A., Guiso, L., and Pistaferri, L. (2018). Portfolio choices, firm shocks, and uninsurable wage risk. *The Review of Economic Studies*, 85(1):437–474.

Fehr, H., Kallweit, M., and Kindermann, F. (2016). Household formation, female labor supply, and savings. *The Scandinavian Journal of Economics*, 118(4):868–911.

Fernández-Blanco, J. (2020). Unemployment risks and intra-household insurance. *Working Paper*.

Fernández-Villaverde, J. and Krueger, D. (2011). Consumption and saving over the life cycle: How important are consumer durables? *Macroeconomic Dynamics*, 15(5):725–770.

Fischer, M. and Khorunzhina, N. (2018). Housing decision under divorce risk. *International Economic Review. Forthcoming*.

Fisher, J. D. and Gervais, M. (2011). Why has home ownership fallen among the young? *International Economic Review*, 52(3):883–912.

Flavin, M. and Yamashita, T. (2011). Owner-occupied housing: life-cycle implications for the household portfolio. *American Economic Review*, 101(3):609–14.

Flood, S., King, M., Rodgers, R., Ruggles, S., and Warren, J. R. (2020). Integrated public use microdata series, current population survey: Version 8.0 [dataset]. minneapolis, mn: Ipums. *Dataset*.

Forbes (2019). The american dream of homeownership is still very much alive. *https://www.forbes.com/sites/fredpeters/2019/04/08/the-american-dream-of-homeownership-is-still-very-much-alive/?sh=3bfa417c3e80*.

Gálvez, J. (2018). Household portfolio choices and nonlinear income risk.

Garcia-Perez, J. and Rendon, S. (2020). Family job search and wealth: the added worker effect revisited. *Quantitative Economics*, 11(4).

Gemici, A. (2011). Family migration and labor market outcomes.

Goetzmann, W. N., Spaenjers, C., and Van Nieuwerburgh, S. (2021). Real and private-value assets. *The Review of Financial Studies*, 34(8):3497–3526.

Goldsmith-Pinkham, P. and Shue, K. (2020). The gender gap in housing returns. *Mimeo*.

Gomes, F. (2020). Portfolio choice over the life cycle: A survey. *Annual Review of Financial Economics*, 12:277–304.

Gomes, F., Haliassos, M., and Ramadorai, T. (2020). Household finance.

Gomes, F. and Michaelides, A. (2005). Optimal life-cycle asset allocation: Understanding the empirical evidence. *The Journal of Finance*, 60(2):869–904.

Goodman, L., Pendall, R., and Zhu, J. (2019). Headship and homeownership–what does the future hold? *Urban Institute Working Paper*.

Goodman, L. S. and Mayer, C. (2018). Homeownership and the american dream. *Journal of Economic Perspectives*, 32(1):31–58.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Griffy, B. S. (2021). Search and the sources of life-cycle inequality. *International Economic Review*.

Guler, B., Guvenen, F., and Violante, G. L. (2012). Joint-search theory: New opportunities and new frictions. *Journal of Monetary Economics*, 59(4):352–369.

Guner, N., Kaygusuz, R., and Ventura, G. (2014). Income taxation of us households: Facts and parametric estimates. *Review of Economic Dynamics*, 17(4):559–581.

Guner, N., Kulikova, Y., and Valladares-Esteban, A. (2020). Does the added worker effect matter? *Working Paper*.

Guvenen, F. (2009). An empirical investigation of labor income processes. *Review of Economic Dynamics*, 12(1):58–79.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of us workers reveal about life-cycle earnings dynamics? *Econometrica*.

Haan, P. and Prowse, V. L. (2017). Optimal social assistance and unemployment insurance in a life-cycle model of family labor supply and savings. *Working Paper*.

Halla, M., Schmieder, J., and Weber, A. (2020). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal: Applied Economics*, 12(4):253–87.

Haurin, D. R., Parcel, T. L., and Haurin, R. J. (2002). Does homeownership affect child outcomes? *Real Estate Economics*, 30(4):635–666.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754.

Heathcote, J., Storesletten, K., and Violante, G. L. (2020). Optimal progressivity with age-dependent taxation. *Journal of Public Economics*, 189:104074.

Heaton, J. and Lucas, D. (2000). Portfolio choice in the presence of background risk. *The Economic Journal*, 110(460):1–26.

Hubener, A., Maurer, R., and Mitchell, O. S. (2015). How family status and social security claiming options shape optimal life cycle portfolios. *The Review of Financial Studies*, 29(4):937–978.

Iskhakov, F., Jørgensen, T. H., Rust, J., and Schjerning, B. (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365.

Jarosch, G. (2015). Searching for job security and the consequences of job loss. *Working Paper*.

Jones, J. B., De Nardi, M., French, E., McGee, R., and Rodgers, R. (2020). Medical spending, bequests, and asset dynamics around the time of death. *NBER WP No 26879*.

Ke, D. (2018). Cross-country differences in household stock market participation: The role of gender norms. In *AEA Papers and Proceedings*, volume 108, pages 159–62.

Khorunzhina, N. and Miller, R. (2019). American dream delayed: shifting determinants of homeownership. *Available at SSRN 3414380*.

Kospentaris, I. (2021). Unobserved heterogeneity and skill loss in a structural model of duration dependence. *Review of Economic Dynamics*, 39:280–303.

Krusell, P., Mukoyama, T., Rogerson, R., and Şahin, A. (2017). Gross worker flows over the business cycle. *The American Economic Review*, 107(11):3447–76.

Krusell, P., Mukoyama, T., and Şahin, A. (2010). Labour-market matching with precautionary savings and aggregate fluctuations. *The Review of Economic Studies*, 77(4):1477–1507.

Lentz, R. (2009). Optimal unemployment insurance in an estimated job search model with savings. *Review of Economic Dynamics*, 12(1):37–57.

Love, D. A. (2010). The effects of marital status and children on savings and portfolio choice. *The Review of Financial Studies*, 23(1):385–432.

Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics*, 3(1, Part 1):11–37.

Lusardi, A. and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1):5–44.

Maloney, T. (1987). Employment constraints and the labor supply of married women: A reexamination of the added worker effect. *Journal of Human Resources*, pages 51–61.

Maloney, T. (1991). Unobserved variables and the elusive added worker effect. *Economica*, pages 173–187.

Mankart, J. and Oikonomou, R. (2016a). Household search and the aggregate labour market. *The Review of Economic Studies*, 84(4):1735–1788.

Mankart, J. and Oikonomou, R. (2016b). The rise of the added worker effect. *Economics Letters*, 143:48–51.

Mankart, J., Oikonomou, R., and Pascucci, F. (2021). The rise of household insurance. *Working Paper*.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, pages 113–126.

Menzio, G. and Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119(3):468–510.

Menzio, G., Telyukova, I. A., and Visschers, L. (2016). Directed search over the life cycle. *Review of Economic Dynamics*, 19:38–62.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413.

Michelacci, C. and Ruffo, H. (2015). Optimal life cycle unemployment insurance. *The American Economic Review*, 105(2):816–59.

Mincer, J. (1978). Family migration decisions. *Journal of Political Economy*, 86(5):749–773.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105(2):385–411.

Mortensen, D. T. (1970). A theory of wage and employment dynamics. *Microeconomic foundations of employment and inflation theory*, pages 167–211.

Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3):397–415.

Mundra, K. and Uwaifo Oyelere, R. (2016). Single and investing: Homeownership trends among the never married.

Nelson, J. A. (1988). Household economies of scale in consumption: theory and evidence. *Econometrica: Journal of the Econometric Society*, pages 1301–1314.

Ortigueira, S. and Siassi, N. (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics*, 60(6):650–666.

Panel Study of Income Dynamics (2021). Produced and distributed by the survey research center, institute for social research, university of michigan, ann arbor, mi (2021). *Dataset*.

Paz-Pardo, G. (2020). Homeownership and portfolio choice over the generations.

Peter, A., Schneider, M., and Piazzesi, M. (2020). Houses and families across countries.

Petrongolo, B. and Pissarides, C. A. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature*, 39(2):390–431.

Piazzesi, M. and Schneider, M. (2016). Housing and macroeconomics. *Handbook of Macroeconomics*, 2:1547–1640.

Poterba, J. M. and Samwick, A. (2001). Household portfolio allocation over the life cycle. In *Aging issues in the United States and Japan*, pages 65–104. University of Chicago Press.

Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296.

Säve-Söderbergh, J. (2012). Self-directed pensions: Gender, risk, and portfolio choices. *The Scandinavian Journal of Economics*, 114(3):705–728.

Sodini, P., Van Nieuwerburgh, S., Vestman, R., and von Lilienfeld-Toal, U. (2021). Identifying the benefits from home ownership: A swedish experiment.

Stephens, M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics*, 20(3):504–537.

Stevenson, B. (2007). The impact of divorce laws on marriage-specific capital. *Journal of Labor Economics*, 25(1):75–94.

Sunden, A. E. and Surette, B. J. (1998). Gender differences in the allocation of assets in retirement savings plans. *The American Economic Review*, 88(2):207–211.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181.

Thörnqvist, T. and Olafsson, A. (2019). Bargaining over risk: The impact of decision power on household portfolios.

Vestman, R. (2019). Limited stock market participation among renters and homeowners. *The Review of Financial Studies*, 32(4):1494–1535.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *The Journal of Finance*, 56(2):433–470.

Vissing-Jorgensen, A. (2002). Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures. *NBER WP No 8884*.

Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8):2295–2332.

Wang, H. (2019). Intra-household risk sharing and job search over the business cycle. *Review of Economic Dynamics*, 34:165–182.

Weinzierl, M. (2011). The surprising power of age-dependent taxes. *The Review of Economic Studies*, 78(4):1490–1518.

Wu, C. and Krueger, D. (2021). Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation. *American Economic Journal: Macroeconomics*, 13(1):79–113.

Yamaguchi, S., Ruiz, C., and Mazzocco, M. (2014). Labor supply, wealth dynamics and marriage decisions.

Yang, F. (2009). Consumption over the life cycle: How different is housing? *Review of Economic Dynamics*, 12(3):423–443.

Yao, R. and Zhang, H. H. (2005). Optimal consumption and portfolio choices with risky housing and borrowing constraints. *The Review of Financial Studies*, 18(1):197–239.

Appendix A

Appendix to Chapter 1

A.1 Data Appendix

A.1.1 The Sample

I work the waves 1989 until 2016 from the Survey of Consumer Finances (SCF) to measure housing and financial choices of households. The SCF is a triennial repeated cross-section analysis sponsored by the Federal Reserve Board. It oversamples asset-rich households, therefore I weigh each observation by the provided survey weights to ensure the representativeness of the US population. For income variables and demographic characteristics I work with data from the Panel Study of Income Dynamics (PSID) spanning from 1989 until 2017.[1] Besides the core sample, the PSID oversamples low-income families (the 'SEO' sample) and immigrant families (the 'immigrant' sample). To make the sample comparable to that from the SCF, I drop these two sub-samples and work with the provided survey weights. In both datasets, I restrict the sample to individuals between 30 and 65 years old. Moreover, I drop the lowest and upper half of a percentile of all financial variables to ensure that results are not driven by individual outliers.

In total, the PSID sample consists of 81,788 individual-year observations that correspond to 2,070 individual single women, 1,589 individual single men and 5,550 individuals living in married couples. The average individual is observed for 5 waves and no individual is observed for more than 15 (biannual)

[1] Because the Survey of Consumer Finances starts in 1989, I restrict my data sample taken from the PSID to the waves from 1989 until 2017. Data were collected annually until 1997 and afterwards every two years.

waves. The data drawn from the SCF (which is a repeated cross-section) includes 39,357 observations, referring to 25,009 individuals in couples, to 4,696 single men and to 7,512 single women.

A.1.2 Supplementary Figures & Tables

In this section, I document additional empirical patterns on housing and financial portfolio composition dynamics of single men, single women and couples in the United States.

Portfolio Choices of Singles by Type

Singles at different ages vary in their marital histories and in their expectations about future marital states which may affect housing choices and portfolio allocation. Figure A.1 plots the life-cycle profiles of median financial assets, conditional home equity and homeownership rates separately for never married and divorced (resp. widowed) singles. Divorced singles are more likely to be homeowners and invest – conditional on owning – less into housing than never married individuals. In contrast, never married singles accumulate more financial assets than divorced individuals.

Figure A.1 Portfolio Allocation of Singles by Type (Data)

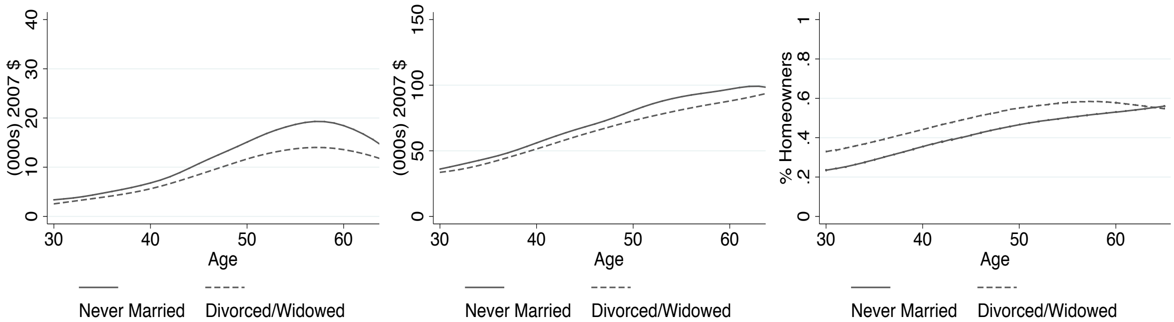

(a) Median Financial Assets **(b)** (Cond.) Housing Equity **(c)** Homeownership Rates

Notes: Figure A.1 plots the life-cycle profiles for median financial assets, housing equity of homeowners and homeownership rates for never married singles as well as for divorced or widowed singles. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

Life-Cycle Patterns of Portfolio Composition

In Section 1.2, I show that the mean house value of owning singles is larger than that of owning couples (per capita). Figure A.2 confirms this (negative) marital gap for median house values, mean home equity and median home equity. Figure A.3a plots median financial assets by family type, confirming that couples accumulate (per capita) more financial assets than singles. When separately considering the extensive and intensive margin of risky asset holdings (Figure A.3), I find that couples are more likely to participate in the stock market. However, conditional on participating, single men accumulate more risky assets than couples, both with regard to the mean and the median of risky asset holdings. Figure A.4 replicates Figure 1.3 but breaks housing equity into mortgages (red bars) and house value (gray bars).

Figure A.2 Housing Choices Across Family Types – Further Specifications

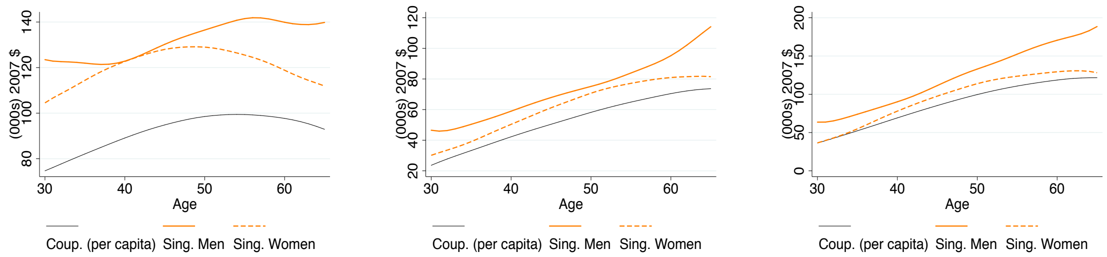

(a) Median House Value **(b)** Median Home Equity **(c)** Mean Home Equity

Notes: Figure A.2 plots the median house value as well as median and mean home equity of owners over the life-cycle. House value is defined as the value of a household's primary residence, irrespective of any mortgage debt. In contrast, home equity refers to the the value of a household's primary residence net of any mortgage debt on this property. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

Children

During their 30s and 40s, more than 80% of couples and more than 60% of single women have children living in their household, whereas only around 20% of single men do.[2] In turn, children may affect households' savings decisions and portfolio allocation. Figure A.5 shows that households with kids are indeed more likely to be homeowners but do not significantly differ from childless households in terms of savings behavior (wealth-to-income ratio) or stock market participation below age 60.[3] The differences beyond 60 arise because households who still have kids in their household at that age are a particular, but small, subsample of the population. Moreover, conditional on household type, differences in homeownership rates by kids disappear for single women and become very small for couples and single men (Figure A.6). Thus, it seems that marital status per se is a more important predictor for portfolio choices than having children. This finding confirms Peter et al. (2020) who show that once they control for being couple or single, children do not explain any additional variation in the housing tenure choice across a sample of European countries.

2 Own calculations from SCF data.

3 "Kids" refers to children that live in the same households or are below 25 and live elsewhere. All Figures look similar when considering only kids who live in the same household or having kids in general.

Figure A.3 Financial Choices Across Family Types – Further Specifications

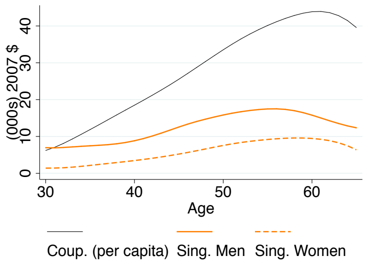

(a) Median Financial Assets

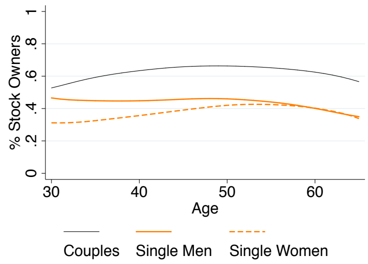

(b) Stock Market Participation

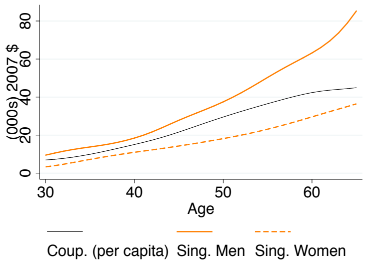

(c) (Cond.) Median Risky Assets

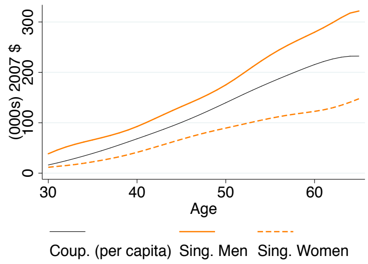

(d) (Cond.) Mean Risky Assets

Notes: Figure A.3 plots the life-cycle profiles of stock market participation rates, median financial assets as well as mean and median risky asset holdings, conditional on participating in the stock market. Financial assets are defined as the sum of safe and risky financial assets. Risky assets contain direct stock holdings, corporate and foreign bonds, the fraction of mutual funds that include the former as well as the fraction of retirement accounts which is invested in stocks. Safe financial assets refer to cash holdings, savings and checking accounts, government bonds and the fraction of mutual funds and retirement accounts which is invested in safe assets. Stock market participation is defined as holding a strictly positive amount of risky assets. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

Housing Expenditure Share by Marital Status & Kids

Figures A.7 and A.8 report housing expenditures shares for singles and couples across the wealth distribution and over the life-cycle. All figures are computed using PSID data. From wave 1999 onwards, households report expenditures on food, transportation, education, health care, children and housing. The latter includes mortgage and loan payments, rent, property taxes, insurance payments, utilities, cable TV, telephone, internet charges, home repairs and home furnishings. I define the housing expenditure share to be the share in overall reported expenditures that a household allocates to the housing category. I find that the housing expenditure of singles is higher than of couples, whereas I do not find any heterogeneity by wealth nor by age (conditional on marital status). Moreover, for singles, the expenditure share on housing is independent on whether or not they live with children in their household. For couples, Figure A.7b displays a higher expenditure share for married households without kids during young ages. Thus, an increase in the number of household members is associated with a decline in the housing expenditure share and this effect is stronger between singles and couples (that is, more adult household members) than across couples with and without children. These findings

Figure A.4 Portfolio Shares by Age – Including Mortgages

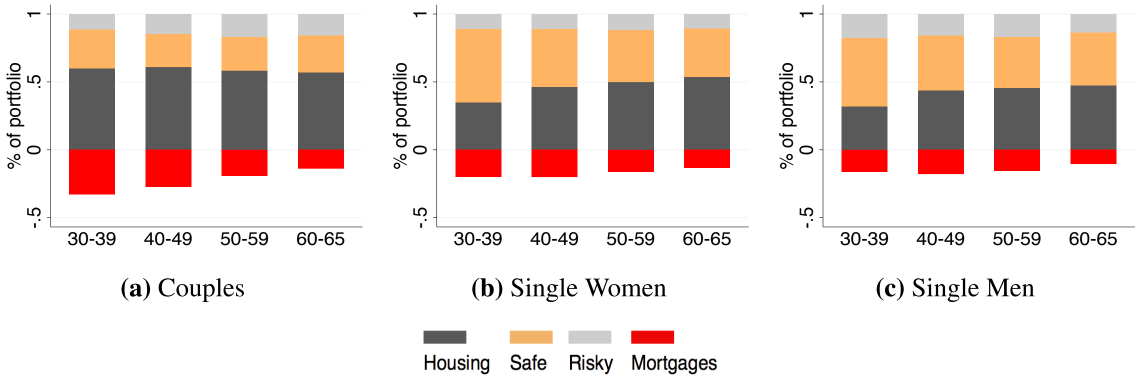

(a) Couples **(b)** Single Women **(c)** Single Men

Notes: Figure A.4 plots the average share of overall wealth invested in housing, mortgages, safe and risky assets by family type and age group. The housing share denotes housing wealth as a fraction of overall wealth (the sum of the house value, safe and risky assets). Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

Figure A.5 Children and Portfolio Composition

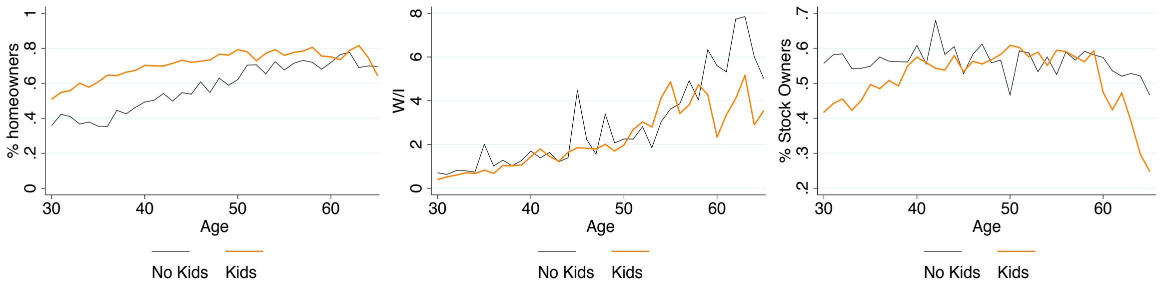

(a) Homeownership Rate **(b)** W/I Ratio **(c)** SMP Rate

Notes: Figure A.5 plots the life-cycle profiles of homeownership rates, wealth-to-income ratios and Stock market participation rate of households with and without kids. "Kids" refer to all children who live in the same household or who are younger than 25 and live elsewhere. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

Figure A.6 Children and Homeownership Rate by Family Type

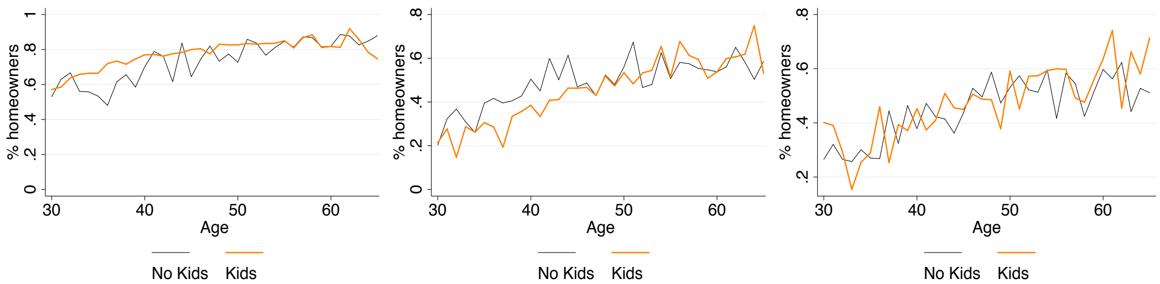

(a) Couples **(b)** Single Women **(c)** Single Men

Notes: Figure A.5 plots the life-cycle of homeownership rates by family type and whether or not the household has kids. "Kids" refer to all children who live in the same household or who are younger than 25 and live elsewhere. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

are in line with Peter et al. (2020) who show that the expenditure share on rent is larger for singles than for couples in Europe.

Figure A.7 Expenditure Shares on Housing across Age

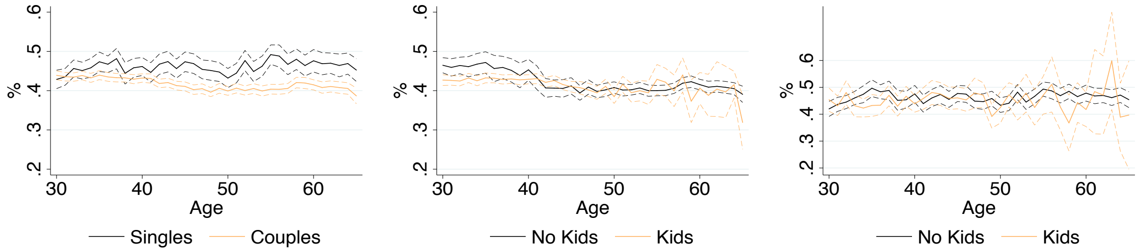

(a) By Marital Status **(b)** By Children – Couples **(c)** By Children – Singles

Notes: Figure A.7 plots the housing expenditure by marital status and by children over the life-cycle. The housing expenditure is defined as expenditures on housing (mortgage and loan payments, rent, property taxes, insurance payments, utilities, cable TV, telephone, internet charges, home repairs and home furnishings) over all reported expenditures categories which include food, transportation, education, health care, children and housing. Data is from the Panel Study of Income Dynamics (PSID), waves 1999-2017.

Figure A.8 Expenditure Shares on Housing across the Wealth Distribution

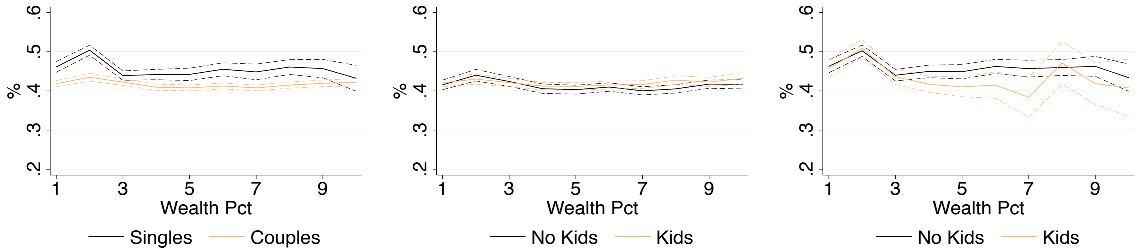

(a) By Marital Status **(b)** By Children – Couples **(c)** By Children – Singles

Notes: Figure A.7 plots the housing expenditure by marital status and by children along the wealth distribution. The housing expenditure is defined as expenditures on housing (mortgage and loan payments, rent, property taxes, insurance payments, utilities, cable TV, telephone, internet charges, home repairs and home furnishings) over all reported expenditures categories which include food, transportation, education, health care, children and housing. Data is from the Panel Study of Income Dynamics (PSID), waves 1999-2017.

Mortgage Characteristics by Family Type

One potential concern in the current analysis is that singles face a different borrowing environment than couples which would render the assumption of homogeneous mortgage premia across all family types unrealistic. To understand the plausibility of this assumption, Table A.1 lists the share of mortgage holders with adjustable loan rates as well as the average mortgage rate across couples, single men and single women in SCF data. Both types of mortgage characteristics do not significantly vary by family type. Additionally, when linearly regressing the mortgage rate on family type while controlling for observable households characteristics (income, mortgage value, age and interview wave), the coefficients for family type turns out to not be statistically significant different from zero.

Table A.1 Mortgage Characteristics by Family Type

	Couples	Singles	
		Men	Women
% with adjustable loan	12.90	12.63	12.24
	(11.73;13.53)	(11.58;12.90)	(12.58;13.22)
Mean mortgage rate in %	6.67	6.58	6.67
	(6.66;6.70)	(6.55;6.66)	(6.66;6.75)

Notes: Table A.1 reports the average mortgage rates and share of households with adjustable rate mortgages by family type. All values are expressed in % and refer to the mortgage that the respective household lists as primary, or "first", mortgage. 95% confidence intervals in parentheses. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

A.2 Robustness Checks – Empirics

A.2.1 Cohabiting Couples

Throughout the benchmark analysis, I drop all couples who cohabit but are not legally married. However, as documented in for example in Adamoupoulou et al. (2021), the share of cohabiting individuals has more than doubled throughout my sample period. Therefore, Figure A.9 and Figure A.10 replicate the main Figures from Section 1.2 when either including cohabiting households in the couples category or in the singles category, respectively (for singles, I allocate cohabiting households to single men if the household head is a man and to single women if the household head is a woman). I do not find any significant differences across specifications. If anything, the homeownership rate of only legally married couples is higher than if I jointly consider married and cohabiting couples. However, and most importantly, it is still substantially higher than that of singles.

Figure A.9 Robustness to Cohabiting Individuals – Couples

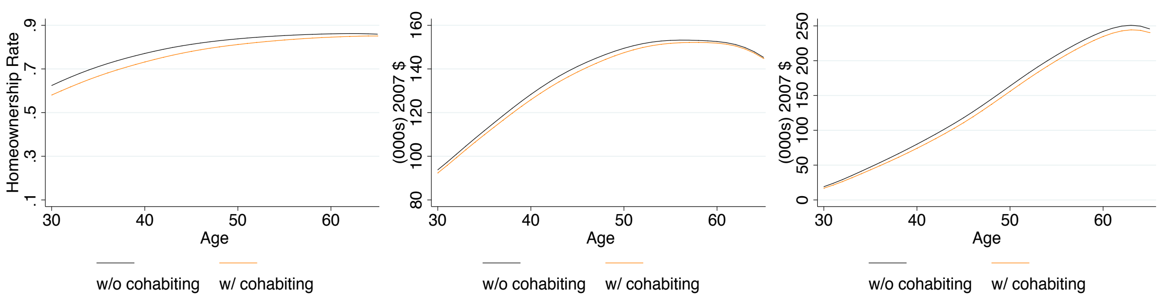

(a) Homeownership Rates **(b)** (Cond.) House Value **(c)** Financial Assets

Notes: Figure A.9 plots homeownership rates, the average house value of owners as well as financial asset accumulation of couples, with and without including cohabiting couples in the couples category (orange and black lines, respectively). Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

Figure A.10 Robustness to Cohabiting Individuals – Singles

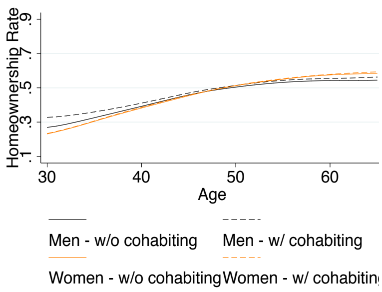

(a) Homeownership Rates

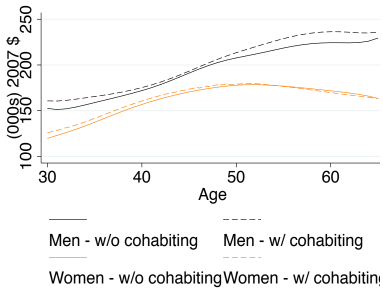

(b) (Cond.) House Value

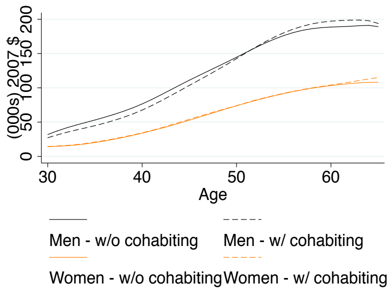

(c) Financial Assets

Notes: Figure A.10 plots homeownership rates, the average house value of owners as well as financial asset accumulation of single men (black lines) and single women (orange lines), with and without including cohabiting couples in the singles category (dashed and solid lines, respectively). Cohabiting couples belong to "single men" if the household reference person is a man and to "single women" if the household reference person is a woman. Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

A.2.2 Cohort Effects

One cannot simultaneously identify age, year and cohort effects because of perfect multi-collinearity. However, Ameriks and Zeldes (2004) show that life-cycle profiles of equity shares look very different depending on whether one imposes either cohorts or year effects to be zero. Throughout my analysis I pool all cohorts who participated in the SCF (resp. PSID). To test the sensitivity of my results to this implicit assumption that cohort effects are zero, Figure A.11 reports the life-cycle profiles of stock market participation rates, homeownership rates, conditional house values and financial assets for individuals who were born between 1945 and 1960. As for the entire sample, I confirm the marital gap in homeownership rates, stock market participants as well as (financial) asset accumulation. Additionally, the conditional house value of singles is higher than that of couples, in line with the benchmark results. Hence, it appears that the reported life-cycle patterns in Section 1.2 are not driven by differences in investment behavior across cohorts.

A.2.3 Excluding Housing Boom and Bust Years

My sample period covers both the housing boom period in the early 2000s as well as the subsequent house price collapse after the financial crisis in 2008. Arguably, both episodes were rather unusual but strongly affected investment patterns. One potential concern is that these episode had heterogeneous effects across family types and hence drive the documented marital gaps in housing choices or in financial portfolio allocation. Figure A.12 reports the life-cycle profiles of homeownership rates, stock market participation rates, conditional house values and financial assets by family type after dropping the years of the housing boom and of the Great Recession (waves 2001, 2004, 2007 and 2010) from

Figure A.11 Portfolio - Robustness to Cohort Effects

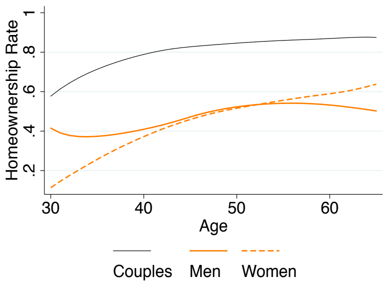

(a) Homeownership Rate

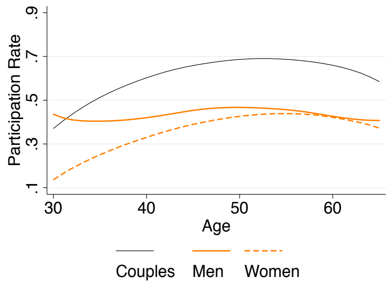

(b) Stock Market Participation Rate

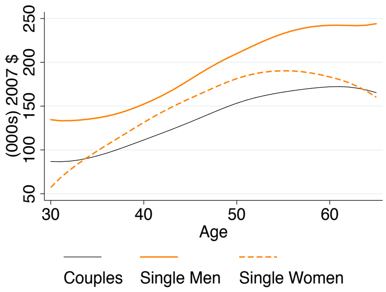

(c) (Cond.) House Value

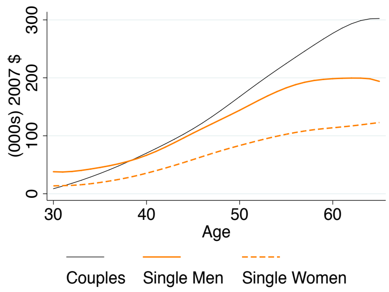

(d) Financial Assets

Notes: Figure A.11 plots homeownership rates, stock market participation rates, the average house value of owners as well as financial asset accumulation by family type on the cohort of individuals born between 1945 and 1960 (in the case of couples, the average birth year across spouses has to be between 1945 and 1960). Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

the sample. I do not find any significant differences in the documented patterns when compared to the benchmark results in Section 1.2.

Figure A.12 Portfolio - Robustness to Boom & Bust Periods

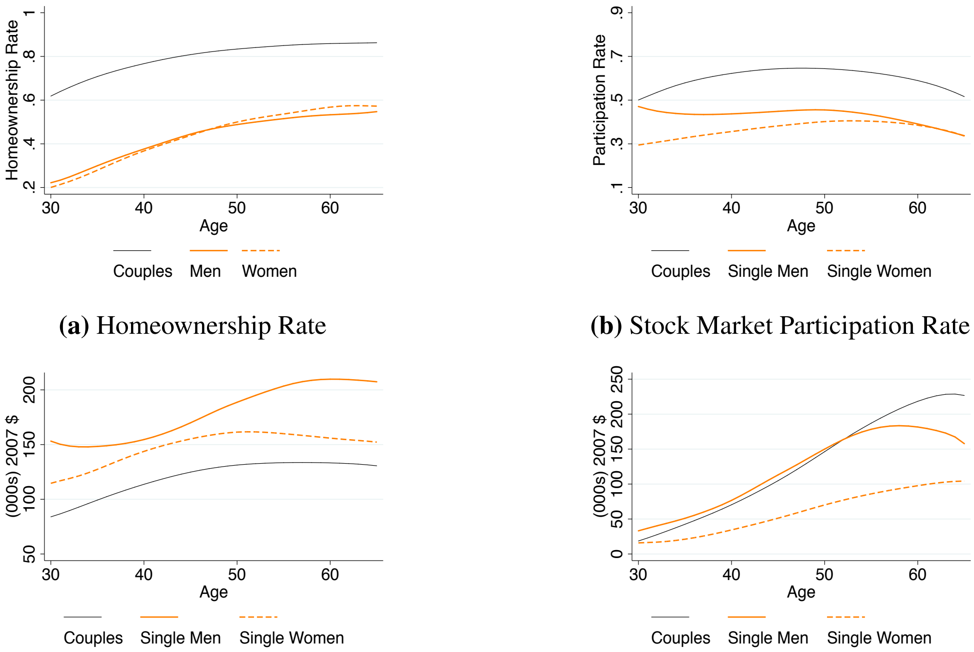

(a) Homeownership Rate

(b) Stock Market Participation Rate

(c) (Cond.) House Value

(d) Financial Assets

Notes: Figure A.12 plots homeownership rates, stock market participation rates, the average house value of owners as well as financial asset accumulation by family type when dropping the waves 2001, 2004, 2007 and 2010 from the sample. Data is from the Survey of Consumer Finances (SCF), waves 1989, 1992, 1995, 1998, 2013 and 2016.

A.3 Model Calibration

A.3.1 First Stage: Income Process – Deterministic Component

I define labor income as annual household income out of labor earnings (including labor income from farms and businesses) and social security benefits converted into 2007 dollars using the CPI-U.[4] I drop households who, according to this measure, report less than \$500 annual income. To estimate the labor income profiles, I follow Borella et al. (2019), split the sample by family type and separately regress the log of income for household i at age j,

$$log(income_{ij}) = \alpha + \beta_1 age_{ij} + \beta_2 age_{ij}^2 + \beta_3 woman_i * age_{ij} + \beta_4 woman_i * age_{ij}^2 + \delta_i + u_{ij}$$

on a fixed effect δ_i, age, age^2 as well as – for singles – their interaction term with a dummy that indicates if the individual is a woman. For singles, to obtain shifters for gender, I regress the sum of the

[4] CPI estimates taken from the US Bureau of Labor Statistics, available under this link [Accessed April 19, 2021].

fixed effect and the residual on a gender dummy:

$$\delta_i + u_{ij} \equiv w_{ij} = \gamma_0 + \gamma_1 woman_i + \epsilon_{ij}$$

The coefficients from these income equations (reported in Table A.2) inform me about the deterministic component of the labor income process which can be split into a constant and an age-specific part.

Table A.2 Regression Coefficients for Income Estimation (Deterministic Component)

	Couples	Singles	
		First Stage	Second Stage
Woman			-1.153***
			(0.0178)
age	0.132***	0.0938***	
	(0.00560)	(0.0116)	
$age^2 * 100$	-0.141***	-0.119***	
	(0.00625)	(0.0123)	
age*woman		0.0198***	
		(0.00539)	
Constant	8.883***	8.616***	0.703***
	(0.122)	(0.272)	(0.0139)
Observations	32,811	13,193	13,193
Number of unique indiv.	5,745	3,467	
R^2	0.045	0.026	0.241

Notes: Estimations are based on (fixed-effect) OLS regressions from PSID Data, waves 1989-2017. Corresponding Figure is Figure 2.4 in the main text. Dependent variable of first stage: Log of annual income (labor income and social security benefits). Dependent variable of second stage: fixed effects plus residual from first stage. *Woman* is a dummy indicating if the individual is woman; Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

A.3.2 First Stage: Income Process – Stochastic Component

I estimate the parameters governing the stochastic part of the income process $\tilde{y}$ with the simulated method of moments, requiring it to match empirical second, third and fourth moments of residual income levels (ϵ_{ij}) in the cross section and for income changes within individuals over time.

Given the functional form of the stochastic income process specified in Section 1.3.3, I need to estimate five parameters per family type. Table A.3 summarizes the estimation results and Table A.4 shows the corresponding data fit. My point estimates imply almost equal persistence across family types. However, singles face larger variances σ_1^2 as well as σ_2^2 and their innovations are less likely to be drawn from the normal distribution with negative mean. The estimated process matches very well the standard deviation and the kurtosis for both income changes and income levels by family type. In addition, it replicates the less negative skewness in income changes for single women, albeit generally predicting too low values for the skewness of income changes. In contrast, it implies slightly too large values for the skewness in the cross sectional dispersion of income realizations when compared to the data.

Table A.3 Estimation Results – Stochastic Income Process

	Couples	Singles	
Parameter		Men	Women
ρ	0.7500	0.7502	0.7505
μ_1	-0.0615	-0.0909	-0.1263
σ_1^2	0.9508	1.4090	2.2888
σ_2^2	0.3141	0.3288	0.4261
$p_{\tilde{y}}$	0.2171	0.1514	0.0425

Notes: Table A.3 presents the estimation results for the stochastic part of the income process by family type, following the parameterization explained in Section 1.3.3.

Table A.4 Data vs. Model – Stochastic Income Process

	Income Levels			Income Changes		
	Couples	Singles		Couples	Singles	
Moment		Men	Women		Men	Women
SD	0.7934	0.9496	0.9524	0.5614	0.6711	0.6737
	0.7834	0.9257	0.9161	0.5665	0.7017	0.6669
Skewness	-0.0969	-0.0932	-0.0412	-0.1629	-0.1565	-0.0611
	-0.1329	-0.1514	-0.2111	-0.1190	-0.1197	-0.0301
Kurtosis	3.9445	3.5814	3.4568	7.5249	9.3101	10.3191
	3.9078	3.6574	3.4522	7.5280	9.3043	10.3260

Notes: Table A.4 compares the second, third and fourth moment for income levels in the cross section as well as for income changes within individuals in the data (gray numbers) with those generated by the simulated income process (black numbers), given the parameter values listed in Table A.3. Data is from the PSID, waves 1989-2017.

A.3.3 First Stage: Marriage and Divorce Probabilities

Figures A.13 plot the life-cycle profiles for divorce and marriage probabilities by productivity realization and, in the case of marriage, by gender. All profiles are obtained by running logit regressions on PSID data whose coefficients are reported in Table A.5.

Figure A.13 Marital Transition Probabilities

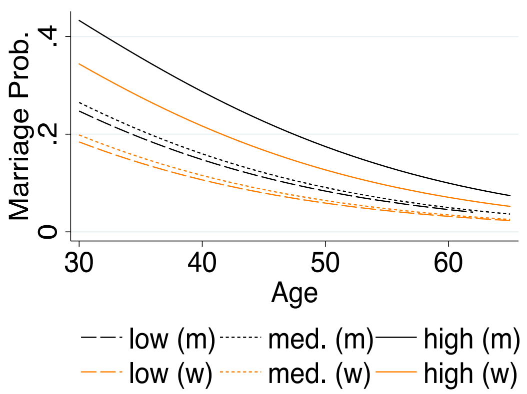

(a) Marriage Probability

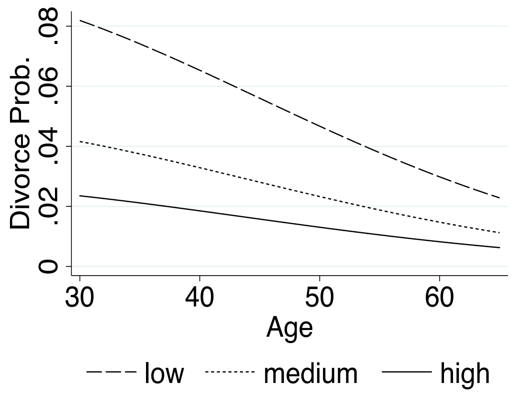

(b) Divorce Probability

Notes: Figure A.13 plots marriage and divorce probabilities by age for individuals with a "low", "medium" and "high" productivity realization, respectively. In Figure A.13a, (m) refers to men and (w) refers to women. Estimates are based on logit regressions whose coefficients are reported in Table A.5. Data is from the Panel Study of Income Dynamics, waves 1989-2017.

Table A.5 Regression Coefficients for Marriage and Divorce Hazards

	(1) Marriage Prob.	(2) Divorce Prob.
Woman	-0.376***	
	(0.0653)	
Age	-0.0627*	0.0143
	(0.0348)	(0.0452)
$age^2 * 100$	0.0018	-0.0554
	(0.0401)	(0.0507)
$\tilde{y}_2$	-0.0852	-0.0221
	(0.324)	(0.533)
$\tilde{y}_3$	-0.385	-0.557
	(0.292)	(0.449)
$\tilde{y}_4$	-0.0891	-0.299
	(0.253)	(0.399)
$\tilde{y}_5$	0.0924	-0.721*
	(0.240)	(0.387)
$\tilde{y}_6$	0.257	-1.056***
	(0.237)	(0.389)
$\tilde{y}_7$	0.199	-1.143***
	(0.243)	(0.411)
$\tilde{y}_8$	0.843***	-1.309**
	(0.285)	(0.641)
$\tilde{y}_9$	0.562	
	(0.527)	
Constant	0.784	-2.348**
	(0.767)	(1.038)
Observations	10,746	27,155

Notes: Estimations are based on Logit regressions from PSID Data, waves 1989-2017. Corresponding Figure is Figure A.13 in the main text. Dependent variable: Likelihood of getting married (resp. divorced) within the next two years, conditional on not being married (resp. being married) today. The age of a couple is the average age across both spouses. *Woman* is a dummy indicating if the individual is woman. $\tilde{y}_x$ is a dummy indicating whether the individual has that productivity realization, with $\tilde{y}_1$ being the base. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

A.4 Further Model Results

A.4.1 Additional Figures on Model Fit

Figure A.14 Portfolio Shares by Age – Data vs. Model (untargeted)

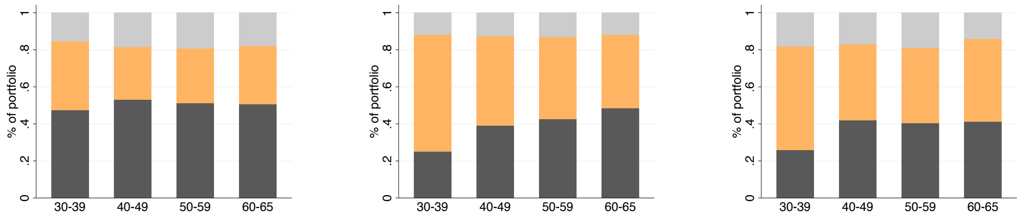

(a) Couples – Data **(b)** Single Women – Data **(c)** Single Men – Data

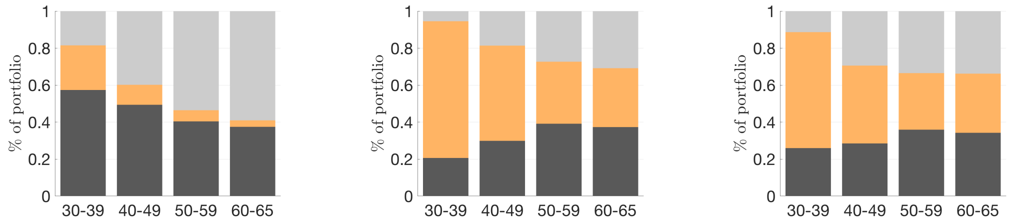

(d) Couples – Model **(e)** Single Women – Model **(f)** Single Men – Model

Notes: Figure A.14 compares portfolio shares by family type from the data (upper panel) with those generated by the model (lower panel). Data is from the Survey of Consumer Finances (SCF), waves 1989-2016.

A.4.2 Additional Policy Functions

Figure A.15 Housing Policy Functions – Single Women

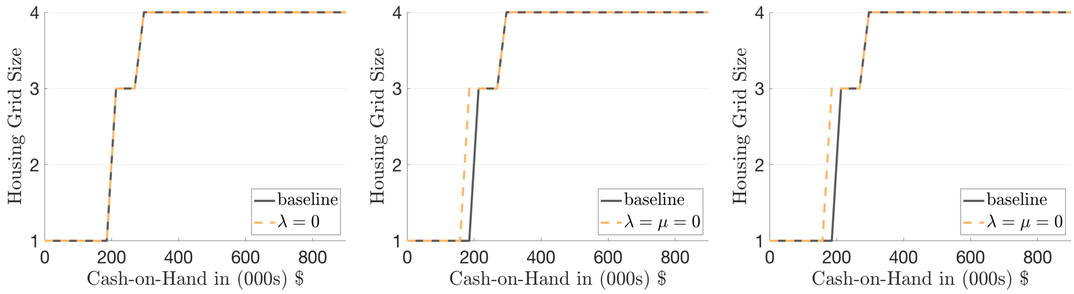

(a) No Divorce **(b)** No Marriage **(c)** No Marital Transitions

Notes: Figure A.15 plots the housing policy functions for single women in the baseline as well as in the counterfactual without divorce ($\lambda = 0$), without marriage ($\mu = 0$) and without any marital transitions ($\lambda = \mu = 0$). All Figures refer to single women of age 30 who rent the smallest house size and who have the medium productivity realization. Figure 1.13 in the main text reports the corresponding policy functions for single men.

Figure A.16 Housing Policy Functions – Couples (Further Factors)

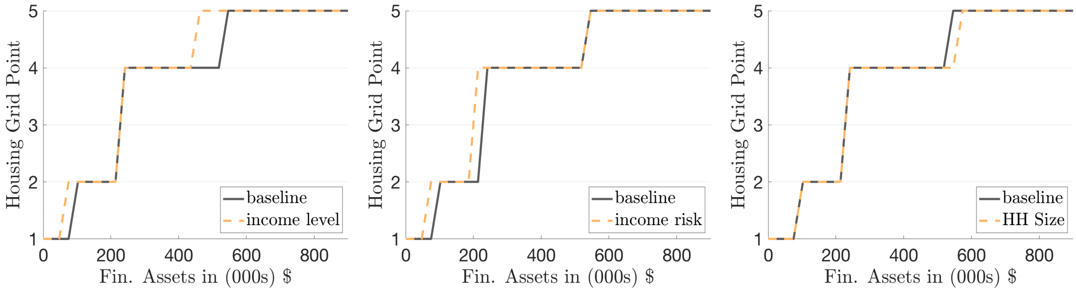

(a) Income Level **(b)** Income Risk **(c)** HH Sizes

Notes: Figure A.16 plots the housing policy functions for couples in the baseline as well as in the counterfactuals that assign singles the income level, income risk, and average number household members of couples, respectively. All Figures refer to couples of age 30 who rent the smallest house size and who have the medium productivity realization.

Figure A.17 Housing Policy Functions – Single Men (Further Factors)

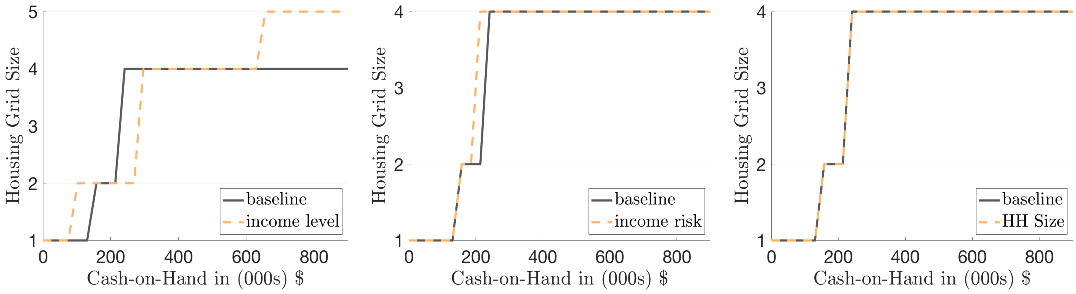

(a) Income Level **(b)** Income Risk **(c)** HH Sizes

Notes: Figure A.17 plots the housing policy functions for single men in the baseline as well as in the counterfactuals that assign singles the income level, income risk, and average number household members of couples, respectively. All Figures refer to single men of age 30 who rent the smallest house size and who have the medium productivity realization.

Figure A.18 Housing Policy Functions – Single Women (Further Factors)

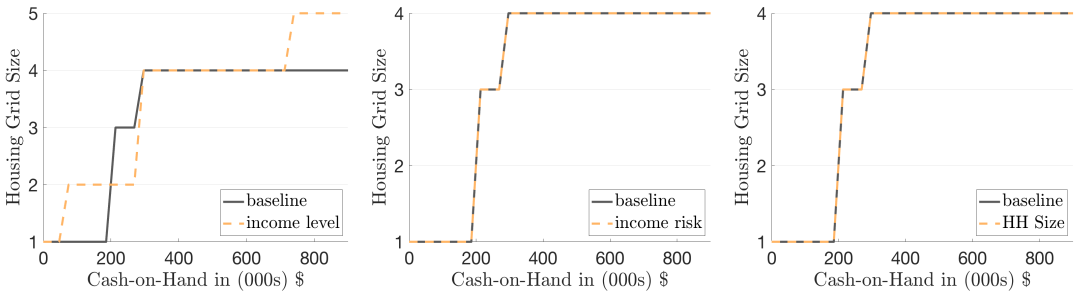

(a) Income Level **(b)** Income Risk **(c)** HH Sizes

Notes: Figure A.18 plots the housing policy functions for single men in the baseline as well as in the counterfactuals that assign singles the income level, income risk, and average number household members of couples, respectively. All Figures refer to single men of age 30 who rent the smallest house size and who have the medium productivity realization.

A.5 Robustness Checks – Model

A.5.1 Moving Frequency

Empirical evidence suggests that singles move more often than couples what shifts their incentives to invest in illiquid housing relative to liquid financial assets (e.g. Blackburn, 2010; Burke and Miller, 2018; Gemici, 2011; Mincer, 1978). Hence, it is possible that higher homeownership rates of couples can be (partially) explained by their lower moving frequency. To test for the importance of this channel, I conduct two exercises. First, Figure A.19 compares moving frequencies by marital status in the data to those generated by the model (without being targeted) and shows that the model replicates these frequencies very well. Second, I introduce an iid moving shock as in Cocco (2005). Each period, households face an exogenous probability to be hit by a moving shock in which case they are forced to move (in the same house size) and have to pay the corresponding adjustment costs. Importantly, the size of the shock is higher for singles to capture their higher incentive to re-locate. Relative to marital transition risk and labor income profiles, the effect of this moving shock on investment choices is quantitatively small and hence, the main results from Sections 1.5 and 1.6 remain unaffected by its introduction.

Figure A.19 Moving Frequencies – Data vs. Baseline Model (untargeted)

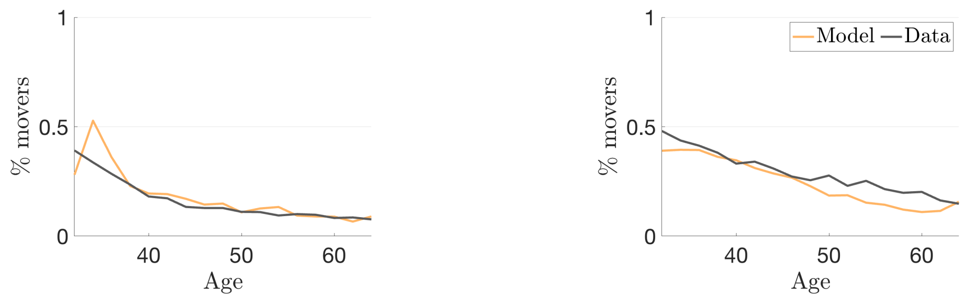

(a) Couples **(b)** Singles

Notes: Figure A.19 plots the moving probabilities by marital status from the data (gray lines) and compares them with model simulations (orange lines). The left graph shows couple households whereas the right graph pools single men and single women, both graphs including owners and renters. Data is from the Panel Study of Income Dynamics (PSID) waves 1997-2017, and refer to the survey question "Did you move since the last interview"?

A.5.2 Housing Grid

In the model, housing is defined over a discrete grid, imposing a minimum threshold into ownership. Hence, it is possible that some low-asset households would find it optimal to buy smaller houses than available in the market and consequently choose to remain renters. Additionally, some households who are located in the right tail of the asset distribution might be constrained by the largest house available, partially explaining the relatively low house values of owning couples in the model. To address this concern, I conduct a robustness exercise in which I increase the housing grid along both directions. In particular, I introduce an additional house that is 1.3 times larger and 1.3 times more expensive than the biggest one previously available. In addition, I decrease the price of the smallest house by 25% and introduce a further housing option whose size and price lies between the (now cheaper) small and medium sized home. Figure A.20 plots the resulting life-cycle profiles and compares them to the baseline framework. The homeownership rate of singles increases substantially when allowing for a smaller house, indicating that the they are (financially) constrained by the lowest housing size. However, the existence of such a threshold can be empirically justified, as in many areas, especially those that singles are most likely to live in (e.g. large cities), even studios are quite expensive and smaller properties are not available in the market. In contrast, the conditional house value of couples does not increase in response to allowing for larger house sizes, indicating that they are not constrained by to few options in the upper part of the housing grid.

Figure A.20 Robustness Exercise – Increasing the Housing Grid

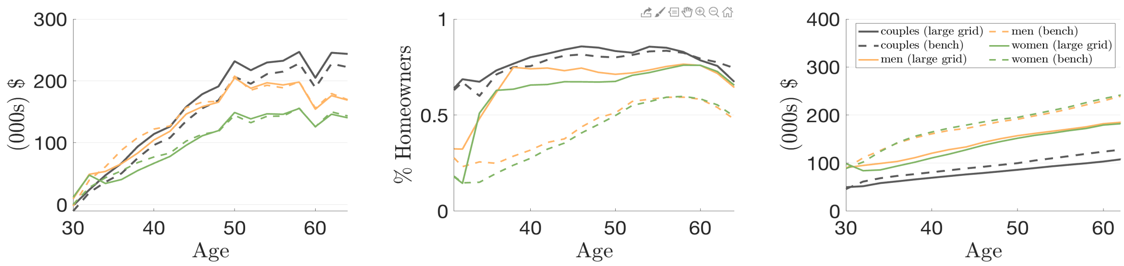

(a) Financial Savings **(b)** Homeownership Rate **(c)** (Cond.) House Value

Notes: Figure A.20 plots the life-cycle profiles of homeownership rates, financial savings and conditional house values across family types in the baseline model (dashed lines) and in version with a larger housing grid (solid lines).

A.5.3 Price Adjustments

I conduct all policy experiments under the implicit assumption that owner-occupied housing supply is fully elastic, i.e. that house prices remain unaffected by the introduction of policy reforms. To test the sensitivity of my results with regard to that simplification, I follow Paz-Pardo (2020) and approximate potential equilibrium effects by re-performing the policy exercises under the assumption that owner-occupied housing supply is characterized by an isoelastic supply function with elasticity $\epsilon = 1.75$, an empirical estimate for the average U.S. metropolitan area by Saiz (2010). To do so, I first compute the housing demand in the baseline model which I define as the number of households i who live in owner-occupied housing, given house prices: $\sum_i H^d(p_H)$. I define this quantity to be the initial housing stock H^s in the economy. Thus, I assume that house prices in the baseline model clear the market: $\sum_i H^d(p_H) = H^s$. Next, I compute the housing demand in each policy counterfactual under baseline prices, that is $\sum_i H^{d'}(p_H)$. Assuming an empirical housing supply elasticity of $\epsilon = 1.75$, the goal is to find the new house prices p'_H, such that:

$$\sum_i H^{d'}(p'_H) = H^{s'}$$

where

$$\epsilon = \frac{\frac{p'_H - p_H}{p_H}}{\frac{H^{s'} - H^s}{H^s}}$$

Hence, I can solve for p'_H by substituting these two equation into one another. To account for different prices across house sizes, I consider the average house price in the economy and assume that all house prices adjust by the same fraction and that they appreciate deterministically as in the benchmark (that is, I do not allow for any segmentation in the housing market). *Panel I* in Table A.6 reports the results. As before, I find that the reduced framework overstates the effectiveness of housing policies and does more so for the case of lowering maintenance costs.

In addition, it is possible that couples who own a house are less likely to separate and hence, divorce rates fall after the introduction of housing policies. In turn, singles may postpone marriage if they are homeowners. Therefore, I re-run the policy exercises from Section 1.6 under the assumption that marriage and divorce rates drop by 20% in response to the housing reforms. *Panel II* in Table A.6 shows that the main results of the paper are robust with respect to these changes in marital transition probabilities.

A.5.4 House Prices

In the model, housing acts as a safe investment. The rationale behind this modeling choice is that I am mostly interested in channels that heterogeneously affect couples and singles and therefore translate into different investment choices. However, all households are equally exposed to house price risk. Moreover, previous literature (Cocco, 2005) has shown that the house price risk does not significantly affect housing demand, because housing primarily serves as a consumption good. In addition, Adelino et al. (2021)

Table A.6 Comparing Policies – Adjusting Prices

	Δ Homeownership Rate	
	$\Phi\downarrow$ $(5\% \to 2\%)$	$\pi\downarrow$ $(1\% \to 0.45\%)$
Panel I: Adjusting House Prices		
Couples	5.16%pts	4.65%pts
Single Men	5.39%pts	6.01%pts
Single Women	3.84%pts	6.25%pts
Aggregate	5.01%pts	5.04%pts
One HH-Type	11.89%pts	13.65%pts
Panel II: Marital transition rates $\downarrow$		
Couples	+6.11%pts	+6.44%pts
Single Men	+4.57%pts	+6.75%pts
Single Women	+4.04%pts	+4.78%pts
Aggregate	+5.63%pts	+6.25%pts

Notes: Table A.6 reports the average increase in homeownership rates in response to lowering housing transaction costs ($\Phi\downarrow$) and lowering maintenance costs ($\pi\downarrow$) under the assumption that housing supply is characterized by an isoelastic supply function with elasticity 1.75 *(Panel I)* and that both marriage and divorce probabilities drop by 20% in response to the introduction of the reforms *(Panel II)*.

document that the majority of US households (71%) perceive housing as a safe investment.[5] Even in 2011, shortly after the financial crisis and the corresponding house price crash, 66% of households considered housing as safe.

A.6 Reduced Framework

The reduced economy is identical to the benchmark economy except that it only contains one generic household type. It can be described by two value functions, one for working age V_W^B and one for retirement V_R^B, respectively:

$$V_W^B(j,a,\mathcal{H},\tilde{y}) = \max_{a_r',a_s',\mathcal{H}',m',c} u(c,s) + \beta \mathbb{E} V_W^B(j+1,a',\mathcal{H}',\tilde{y}')$$

$$a_r' + a_s' - m' + c = a + p_h\mathcal{H} - p_h\mathcal{H}' - \mathbb{1}_{\mathcal{H}'\neq\mathcal{H}}\Phi(\mathcal{H},\mathcal{H}') - \mathbb{1}_{a_r'>0}S^F - \mathbb{1}_{\mathcal{H}=R}\alpha p_H H_1 - \mathbb{1}_{\mathcal{H}\neq R}\pi\mathcal{H}$$

$$m' \leq \zeta_h p_h \mathcal{H}' \qquad a = \sum_{l=r,s}(1+(1-\tau_k)r_l)a_l - (1+r_m)m + \mathbb{Y}[y(j,\tilde{y}),m]$$

5 These numbers are based on a nationally representative housing survey from Fannie Mae of more than 50,000 households between 2010 and 2016.

$$V_R^B(j,a,\mathcal{H},\hat{y}) = \max_{a'_s,a'_r,\mathcal{H}',m',c} u(c,s) + \beta\psi_j \mathbb{E}V_R^B(j+1,a',\mathcal{H}',\hat{y}) + \beta(1-\psi_j)L\frac{(\xi+a'+\mathcal{H}')^{1-\gamma}}{1-\gamma}$$

$$a'_r + a'_s - m' + c = a + p_h\mathcal{H} - p_h\mathcal{H}' - \mathbb{1}_{\mathcal{H}'\neq\mathcal{H}}\Phi(\mathcal{H},\mathcal{H}') - \mathbb{1}_{a'_r>0}S^F - \mathbb{1}_{\mathcal{H}=R}\alpha_R p_H H_1 - \mathbb{1}_{\mathcal{H}\neq R}\pi\mathcal{H}$$

$$m' \leq \zeta_h p_h \mathcal{H}' \qquad m_J = 0 \qquad a = \sum_{l=r,s}(1+(1-\tau_k)r_l)a_l - (1+r_m)m + \mathbb{Y}(pen(\hat{y}),m)$$

To calibrate the reduced framework, I re-estimate all model elements which are allowed to vary by family type in the benchmark for the pooled sample: income profiles (both in terms of level and risk), average household sizes and survival probabilities. Moreover, I use the tax parameters for the entire population provided in Guner et al. (2014). All remaining parameters (including preference values) are held constant when compared to the benchmark. Table A.7 shows the data fit for key moments in the reduced framework.

Table A.7 Model Fit – One HH-Type Economy

	Data	Model
W/I at 45	1.24	1.55
Mean SMP at 45	48%	51%
homeownership rate at 45	61%	56%
Mean (cond.) house value at 45	\$257,634	\$201,664

Notes: Table A.7 reports the model fit for the reduced framework with one generic household type. Data values refer to the pooled sample in the Survey of Consumer Finances (SCF), waves 1989-2016.

A.7 Policy Exercises – Wealth Accumulation

Table A.8 splits the increase in average household net worth in response to the reforms discussed in Section 1.6 into changes in housing wealth and into changes in financial wealth, corresponding to Table 1.6 in the main text. Both housing policies increase housing investment at the expense of financial savings, while decreasing S^F results in households substituting away from housing towards risky assets, slightly decreasing aggregate housing wealth.

-

Table A.8 Comparing Policies – Housing and Financial Wealth

	$\Phi \downarrow$ $(5\% \rightarrow 2\%)$	$\pi \downarrow$ $(1\% \rightarrow 0.35\%)$	$S^F \downarrow$ $(\$1,275 \rightarrow \$713)$
		Δ Housing Wealth in \$	
Panel I: Bench			
Couples	20,999	14,803	306
Single Men	7,092	5,444	-4,344
Single Women	10,066	10,655	-505
Aggregate	16,935	12,634	-552
Panel II: One HH-Type	18,242	20,645	-1,383
		Δ Financial Wealth in \$	
Panel I: Bench			
Couples	-12,572	-7,935	9,736
Single Men	-10,793	-7,015	10,340
Single Women	-8.720	-6,480	6,673
Aggregate	-11,619	-7,537	9,288
Panel II: One HH-Type	-15,298	-13,630	7,152

Notes: Table A.8 reports changes in housing wealth and in financial wealth in response to lowering housing adjustment costs ($\Phi \downarrow$), lowering housing maintenance costs ($\pi \downarrow$) and lowering stock market participation costs ($S^F \downarrow$) in the benchmark economy *(Panel I)* as well as in the reduced economy with one household type *(Panel II)*.

Appendix B

Appendix to Chapter 2

B.1 Supplementary Figures

One concern in my study could be that the gender investment gap mainly arises through asset holdings in retirement accounts. If single men are more likely to hold retirement accounts (e.g. because of their job types or employment histories) than single women, and if individuals, regardless of gender, tend to invest retirement savings more risky than other types of wealth, the gender investment gap would reflect gender heterogeneity in the labor market rather than in investment choices. Therefore, Figure B.1 plots the life-cycle profiles of equity shares, stock market participation rates and conditional risky shares for singles and couples based on a tighter definition of financial assets that excludes assets held through retirement accounts. It shows that, compared to the baseline, the gender gap in equity shares slightly increases (Figure B.1a), alleviating concerns that empirical investment differences across single men and single women are mainly driven through forms of savings that are linked to certain types of jobs.

Figure B.1 Life Cycle Pattern of Household Finances – Excluding Retirement Accounts

(a) Equity Share **(b)** Participation Rate **(c)** Conditional Risky Share

Notes: Figure B.1 plots the life-cycle profiles of the equity share, stock market participation rates and conditional risky shares for singles and couples. The sample consists of individuals born between 1945 and 1960 in the waves 1989 until 2016 of the Survey of Consumer Finances (SCF). Risky assets are defined as direct stock holdings, corporate and foreign bonds as well as the fraction of mutual funds that include the former. In contrast to Figure 2.1, financial assets do not include wealth held through retirement accounts.

B.2 Regression Coefficients and Marginal Effects

Table B.1 Regression Coefficients & Marginal Effects – Participation Rates of Singles

	(1) SMP	(2) SMP	(3) SMP
single women	-0.278***	-0.130***	-0.0377**
	(0.0161)	(0.0158)	(0.0159)
single woman * age	0.00470***	0.00206***	0.000588*
	(0.000309)	(0.000291)	(0.000305)
age	-0.0287**	-0.0413**	-0.0898***
	(0.0139)	(0.0184)	(0.0222)
$age^2 * 100$	0.114***	0.106***	0.204***
	(0.0308)	(0.0387)	(0.0477)
$age^3 * 10000$	-0.112***	-0.0841***	-0.153***
	(0.0222)	(0.0266)	(0.0332)
High education		0.269***	0.114***
		(0.00402)	(0.00487)
No. of HH members		-0.0305***	-0.0187***
		(0.00197)	(0.00159)
Income		0.0290***	0.0175***
		(0.000552)	(0.000547)
Safe assets			0.0559***
			(0.000593)
Constant	0.464**	0.415	0.974***
	(0.201)	(0.273)	(0.328)
Observations	4,737	4,735	4,735
R^2	0.014	0.126	0.293
Year FE	No	Yes	Yes
ME for women at age 30	0.00452	-0.00639***	-0.00244
	(0.00338)	(0.00257)	(0.00319)
ME for women at mean age (50)	0.197***	0 .0777***	0.0216
	(0.0154)	(0.0138)	(0.0151)
ME for women at age 65	0.334***	0.138***	0.0387
	(0.0243)	(0.0222)	(0.024)

Notes: Estimations are based on OLS on the sample of individuals that live in households with no spouse present. Source: SCF waves 1989 until 2016; individuals born between 1945 and 1960. SMP = Stock Market Participation. *single woman* is a dummy indicating that the household head is a women. *high education* is a dummy equal to one if the household head has more than 12 years of education. *safe assets* refers to safe liquid assets. "ME" indicates the marginal effect of being a women at the respective age. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table B.2 Regression Coefficients & Marginal Effects – Conditional Risky Share of Singles

	(1) Cond. Share	(2) Cond. Share	(3) Cond. Share
single women	0.0515	0.0620*	0.0177
	(0.0320)	(0.0332)	(0.0252)
single woman * age	-0.00234***	-0.00245***	-0.00175***
	(0.000628)	(0.000647)	(0.000511)
age	0.172***	0.0923	0.0499
	(0.0593)	(0.0602)	(0.0571)
$age^2 * 100$	-0.302**	-0.169	-0.0740
	(0.124)	(0.128)	(0.120)
$age^3 * 10000$	0.170**	0.101	0.0374
	(0.0850)	(0.0880)	(0.0828)
No. of HH members		-0.00901**	-0.0176***
		(0.00371)	(0.00345)
Income		0.00304***	0.00543***
		(0.000841)	(0.000754)
Safe assets			-0.0604***
			(0.000671)
Constant	-2.635***	-1.244	-0.0939
	(0.918)	(0.919)	(0.880)
Observations	2,173	2,173	2,173
R^2	0.034	0.054	0.223
Year FE	No	Yes	Yes
ME for women at age 30	-0.0887***	-0.085***	-0.0871***
	(0.00649)	(0.00632)	(0.00672)
ME for women at mean age (50)	-0.187	-0.188***	-0.161***
	(0.0324)	(0.0331)	(0.0275)
ME for women at age 65	-0.252***	-0.256***	-0.209***
	(0.0498)	(0.0511)	(0.0417)

Notes: Estimations are based on OLS on the sample of individuals that live in households with no spouse present. Source: SCF waves 1989 until 2016; individuals born between 1945 and 1960. Cond. Share = Risky Share conditional on Participation. *single woman* is a dummy indicating that the household head is a women. *high education* is a dummy equal to one if the household head has more than 12 years of education. *safe assets* refers to safe liquid assets. "ME" indicates the marginal effect of being a women at the respective age. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

B.3 Model – First Stage Estimation

B.3.1 Income Process Estimation – Stochastic Component

I estimate the stochastic component of the income process by the minimum distance estimator as in Guvenen (2009). In particular, I assume the unexplainable part of the income process (that is, the residual term ϵ_{it} from the income equation) to follow a persistent-transitory process:

$$\tilde{y}_j = z_j + \epsilon_{\tilde{y}}$$

$$z_{j+1} = \rho_z z_j + \nu_z$$

A persistent-transitory process requires identification of three parameters: The persistence parameter ρ_z, the variance of the persistent shock $\sigma_{\epsilon^2_{\tilde{y}}}$ and the variance of the transitory shock $\sigma^2_{\nu_z}$ which can be identified by the following moments:

$$\frac{cov(\tilde{y}_j, \tilde{y}_{j-2})}{cov(\tilde{y}_{j-1}, \tilde{y}_{j-2})} = \frac{\rho_z^2 var(z_{j-2})}{\rho_z var(z_{j-2})} = \rho_z$$

$$var(\tilde{y}_{j-1}) - \frac{cov(\tilde{y}_j, \tilde{y}_{j-1})}{\rho_z} = var(z_{j-1}) + \sigma_{\epsilon_{\tilde{y}}} - var(z_{j-1}) = \sigma_{\epsilon_{\tilde{y}}}$$

$$var(\tilde{y}_{j-1}) - cov(\tilde{y}_j, \tilde{y}_{j-2}) - \sigma_{\epsilon_{\tilde{y}}} = \rho_z^2 var(z_{j-2}) + \sigma_{\nu_z} + \sigma_{\epsilon_{\tilde{y}}} - \rho_z^2 var(z_{j-2}) - \sigma_{\epsilon_{\tilde{y}}} = \sigma_{\nu_z}$$

I recover the parameters that minimize the distance between the covariance-variance matrices of the income process in the data and their theoretical counterparts under the assumption that $Var(z_{-1}) = 0$. However, the PSID only collects data every two years after 1997 while the model is written in annual frequency. To account for this inconsistency, I linearly interpolate income for individuals that I observe in two consecutive waves for the missing year in which no PSID data was collected. I run four different estimations for married men, married women, single men and single women. Table 2.2 in the main text displays the results.

B.3.2 Income Process – Deterministic Component

Table B.3 Regression Coefficients for Income Estimation (Deterministic Component)

	(1) First Stage	(2) Second Stage	(3) First Stage	(4) Second Stage
	Couples		Singles	
high educ.		0.660***		0.618***
		(0.0248)		(0.0398)
Woman		-3.432***		-0.573***
		(0.0476)		(0.0428)
Woman*high educ.		-0.484***		0.204***
		(0.0593)		(0.0527)
age	0.148***		0.0553**	
	(0.0251)		(0.0224)	
$age^2 * 100$	-0.155***		-0.0672***	
	(0.0260)		(0.0225)	
age*woman	0.0449***		0.00580	
	(0.00582)		(0.00539)	
Constant	6.144***	1.390***	9.841***	-0.140***
	(0.610)	(0.0203)	(0.557)	(0.0320)
Observations	34,280	34,280	5,722	5,722
Number of unique indiv.	3,284		892	
R^2	0.017	0.348	0.008	0.169

Notes: Estimations are based on (fixed-effect) OLS regressions from PSID Data, waves 1989-2017 on individuals born between 1945 and 1960. Corresponding Figure is Figure 2.4 in the main text. Dependent variable of first stage: Log of annual income (labor income, social security income and transfers) of the household head. In years where social security (transfer) income is not available separately by head and spouse, I use combined social security (transfer) income and assign it 50-50 to both spouses. For singles, I add labor income, social security benefits and transfers from other household members. For couples, I again split the income from other household members 50-50 between spouses. Dependent variable of second stage: fixed effects plus residual from first stage. *high educ.* is a dummy equal to one if the individual has more than 12 years of schooling; *Woman* is a dummy indicating if the individual is woman; Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

B.3.3 Marriage and Divorce Probabilities

Table B.4 Regression Coefficients for Marriage and Divorce Hazards

	(1) Marriage Prob.	(2) Divorce Prob.
Woman	-0.466***	
	(0.102)	
Age	0.134**	0.0705
	(0.0611)	(0.0513)
$age^2 * 100$	-0.201***	-0.0897
	(0.0675)	(0.0548)
$\mathbb{1} > 1997$	0.218	0.325***
	(0.135)	(0.109)
High Educ. (Head)	-0.112	-0.476***
	(0.101)	(0.0922)
High Educ. (Spouse)		-0.221**
		(0.0903)
Constant	-4.190***	-5.060***
	(1.363)	(1.179)
Observations	7,489	38,104

Notes: Estimations are based on Logit regressions from PSID Data, waves 1989-2017 on individuals born between 1945 and 1960. Corresponding Figure is Figure A.13 in the main text. Dependent variable: Likelihood of getting married (resp. divorced) within the next year, conditional on not being married (resp. being married) today. The age of a couple is the average age of both spouses. For education within couple, head refers to the husband and spouse refers to the wife. In contrast, singles are always labeled as head. *High Educ.* is a dummy equal to one if the individual has more than 12 years of schooling; *Woman* is a dummy indicating if the individual is woman; $\mathbb{1} > 1997$ indicates observations that were interviewd after 1997 to account for the changing frequency of the PSID. Robust standard errors in parentheses, *** p$<$0.01, ** p$<$0.05, * p$<$0.1

Appendix C

Appendix to Chapter 3

C.1 Empirical Robustness Exercises

C.1.1 Employed and Unemployed Spouses

Table C.1 Joint Labor Market Transitions (Full Sample): Spouse Unemployed

	Primary earner transition		
	EE	EU	EN
Cond. prob. of spousal UE transition	25.29%	26.27%	34.11%
Cond. prob. of spousal UU transition	61.97%	63.33%	46.01%
Cond. prob. of spousal UN transition	12.74%	10.41%	19.87%

Notes: This table shows the probability of a spousal transition from unemployment conditional on primary earner transitions for the entire population.

Table C.2 Joint Labor Market Transitions (Full Sample): Spouse Employed

	Primary earner transition		
	EE	EU	EN
Cond. prob. of spousal EE transition	97.61%	91.49%	88.84%
Cond. prob. of spousal EU transition	0.77%	5.78%	1.25%
Cond. prob. of spousal EN transition	1.62%	2.72%	9.92%

Notes: This table shows the probability of a spousal transition from employment conditional on primary earner transitions for the entire population.

C.1.2 Education

Table C.3 Joint Labor Market Transitions by Spousal Education

	Primary earner transition		
	EE	EU	EN
I. Spouse College Degree (All):			
Cond. prob. of spousal NE transition	6.91%	11.40%	20.88%
Cond. prob. of spousal NU transition	1.59%	6.43%	1.04%
Cond. prob. of spousal NN transition	91.50%	82.18%	78.08%
II. Spouse No College Degree (All):			
Cond. prob. of spousal NE transition	5.55%	7.20%	15.08%
Cond. prob. of spousal NU transition	1.65%	5.34%	1.45%
Cond. prob. of spousal NN transition	92.81%	87.46%	83.47%
III. Spouse College Degree (Young):			
Cond. prob. of spousal NE transition	7.31%	13.25%	33.25%
Cond. prob. of spousal NU transition	1.70%	7.22%	1.29%
Cond. prob. of spousal NN transition	90.99%	79.53%	65.46%
IV. Spouse College Degree (Old):			
Cond. prob. of spousal NE transition	6.04%	7.72%	11.81%
Cond. prob. of spousal NU transition	1.35%	4.87%	0.86%
Cond. prob. of spousal NN transition	92.61%	87.41%	87.33%
V. Spouse No College Degree (Young):			
Cond. prob. of spousal NE transition	6.30%	8.34%	21.76%
Cond. prob. of spousal NU transition	2.01%	6.28%	2.21%
Cond. prob. of spousal NN transition	91.69%	85.37%	76.03%
VI. Spouse No College Degree (Old):			
Cond. prob. of spousal NE transition	4.19%	4.20%	9.41%
Cond. prob. of spousal NU transition	0.99%	2.83%	0.80%
Cond. prob. of spousal NN transition	94.82%	92.97%	89.79%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions by education of the spouse.

C.1.3 Cohort Effects

Table C.4 Joint Labor Market Transitions

	Primary earner transition		
	EE	EU	EN
I. Spouse is a Man (Young) :			
Cond. prob. of spousal NE transition	13.54%	14.07%	44.10%
Cond. prob. of spousal NU transition	6.19%	11.69%	2.59%
Cond. prob. of spousal NN transition	80.27%	74.24%	53.31%
II. Spouse is a Man (Old):			
Cond. prob. of spousal NE transition	4.50%	4.59%	10.36%
Cond. prob. of spousal NU transition	1.13%	3.23%	0.63%
Cond. prob. of spousal NN transition	94.37%	92.18 %	89.01%
III. Spouse born between 1960-70 (Young):			
Cond. prob. of spousal NE transition	6.98%	8.62%	21.67%
Cond. prob. of spousal NU transition	1.89%	6.70%	2.42%
Cond. prob. of spousal NN transition	91.13%	84.68%	75.92%
IV. Spouse born between 1960-70 (Old)			
Cond. prob. of spousal NE transition	4.28%	2.94%	12.86%
Cond. prob. of spousal NU transition	1.11%	3.68%	1.04%
Cond. prob. of spousal NN transition	94.61%	93.38%	86.10%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions by gender and cohort.

C.1.4 Children

Table C.5 Joint Labor Market Transitions (< Age 40)

	Primary earner transition		
	EE	EU	EN
I. Have Children:			
Cond. prob. of spousal NE transition	6.26%	8.71%	28.30%
Cond. prob. of spousal NU transition	1.75%	6.65%	2.31%
Cond. prob. of spousal NN transition	91.98%	84.64%	69.40%
II. No Children:			
Cond. prob. of spousal NE transition	9.68%	12.68%	23.69%
Cond. prob. of spousal NU transition	3.40%	8.54%	1.59%
Cond. prob. of spousal NN transition	86.91%	78.78%	74.72%
III. Have Children below 5:			
Cond. prob. of spousal NE transition	5.63%	8.55%	30.09%
Cond. prob. of spousal NU transition	1.47%	6.14%	1.96%
Cond. prob. of spousal NN transition	92.90%	85.31%	67.95%
IV. No Children below 5:			
Cond. prob. of spousal NE transition	8.08%	9.95%	24.82%
Cond. prob. of spousal NU transition	2.60%	7.80%	2.35%
Cond. prob. of spousal NN transition	89.32%	82.24%	72.82%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions by presence of children in the household.

C.1.5 Reasons for Non-Participation

Table C.6 Joint Labor Market Transitions

	Primary earner transition		
	EE	EU	EN
I. Excluding Retirement (Young):			
Cond. prob. of spousal NE transition	6.66%	9.32%	27.13%
Cond. prob. of spousal NU transition	2.00%	6.91%	2.06%
Cond. prob. of spousal NN transition	91.33%	83.77%	70.81%
II. Excluding Retirement (Old):			
Cond. prob. of spousal NE transition	4.95%	4.15%	11.45%
Cond. prob. of spousal NU transition	1.18%	3.33%	1.00%
Cond. prob. of spousal NN transition	93.87%	92.52%	87.54%
III. Excluding Disabled/Ill (Young):			
Cond. prob. of spousal NE transition	6.55%	9.34%	27.02%
Cond. prob. of spousal NU transition	1.96%	6.94%	2.01%
Cond. prob. of spousal NN transition	91.49%	83.72%	70.97 %
IV. Excluding Disabled/Ill (Old):			
Cond. prob. of spousal NE transition	4.17%	3.42%	8.53%
Cond. prob. of spousal NU transition	0.88%	2.77%	0.50%
Cond. prob. of spousal NN transition	94.95%	93.81%	90.97%
V. Excluding Retired and Disabled/Ill (Young):			
Cond. prob. of spousal NE transition	6.55%	9.36%	27.23%
Cond. prob. of spousal NU transition	1.97%	6.96%	2.05%
Cond. prob. of spousal NN transition	91.48%	83.68%	70.72%
VI. Excluding Retired and Disabled/Ill (Old):			
Cond. prob. of spousal NE transition	4.74%	3.62%	11.20%
Cond. prob. of spousal NU transition	1.16%	3.40%	0.89%
Cond. prob. of spousal NN transition	94.11%	92.99%	87.91%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions by reasons for non-participation.

C.1.6 Business Cycle

Table C.7 Joint Labor Market Transitions

	Primary earner transition		
	EE	EU	EN
NBER Recession, Young			
Cond. prob. of spousal NE transition	6.48%	7.74%	22.38%
Cond. prob. of spousal NU transition	1.98%	8.73%	0.99%
Cond. prob. of spousal NN transition	91.55%	83.53%	76.63%
NBER Recession, Old			
Cond. prob. of spousal NE transition	4.14%	5.43%	7.71%
Cond. prob. of spousal NU transition	0.83%	2.76%	0.68%
Cond. prob. of spousal NN transition	95.03%	91.81%	91.61%
No NBER Recession, Young			
Cond. prob. of spousal NE transition	6.68%	9.53%	27.45%
Cond. prob. of spousal NU transition	2.00%	6.63%	2.14%
Cond. prob. of spousal NN transition	91.31%	83.85%	70.41%
No NBER Recession, Old			
Cond. prob. of spousal NE transition	4.30%	3.46%	8.80%
Cond. prob. of spousal NU transition	0.91%	2.75%	0.54%
Cond. prob. of spousal NN transition	94.79%	93.79%	90.66%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions by state of the business cycle.

C.1.7 Income

Table C.8 Joint Labor Market Transitions by Past Income

	Primary earner transition		
	EE	EU	EN
I. Low Income (All):			
Cond. prob. of spousal NE transition	5.57%	7.41%	15.89%
Cond. prob. of spousal NU transition	2.98%	5.81%	1.64%
Cond. prob. of spousal NN transition	92.45%	86.79%	82.48%
II. High Income (All):			
Cond. prob. of spousal NE transition	5.91%	8.93%	20.71%
Cond. prob. of spousal NU transition	1.14%	4.75%	0.73%
Cond. prob. of spousal NN transition	92.95%	86.32%	78.55%
III. Low Income (Young):			
Cond. prob. of spousal NE transition	6.22%	8.66%	23.30%
Cond. prob. of spousal NU transition	2.37%	7.48%	2.38%
Cond. prob. of spousal NN transition	91.41%	83.85%	74.32%
IV. Low Income (Old):			
Cond. prob. of spousal NE transition	3.66%	3.52%	8.11%
Cond. prob. of spousal NU transition	0.95%	2.41%	0.66%
Cond. prob. of spousal NN transition	95.39%	94.08%	91.24%
V. High Income (Young):			
Cond. prob. of spousal NE transition	7.24%	7.13%	40.17%
Cond. prob. of spousal NU transition	1.19%	4.18%	0.12%
Cond. prob. of spousal NN transition	91.57%	88.69%	59.71%
VI. High Income (Old):			
Cond. prob. of spousal NE transition	4.76%	3.66%	11.121%
Cond. prob. of spousal NU transition	0.90%	2.84%	0.49%
Cond. prob. of spousal NN transition	94.34%	93.50%	88.30%

Notes: This table shows the probability of a spousal transition from out of the labor force conditional on primary earner transitions by income.

C.1.8 Dynamics Response for Other Age Groups

Figure C.1 Δ Pr(Spouse enters LF) <u>this month</u>

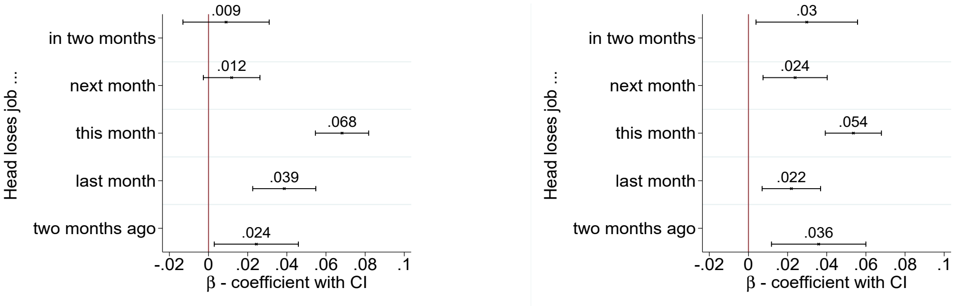

(a) Age 36 to 45 **(b)** Age 46 to 55

Notes: Figure C.1 shows the change in probability that a non-participating spouse enters the labor force (either as unemployed or as employed) if the household head loses/lost the job in two months, next month, this month, last month or two months ago, respectively, relative to the baseline in which the household head remains employed. The sample includes couples in which one spouse is working and one spouse is out of the labor force between age 36 and 45 (Figure C.1a) and between age 46 and 55 (Figure C.1b) from the Current Population Survey (CPS), waves 1994 until 2020. Age refers to the non-participating spouse. The regression producing the coefficients is Equation 3.1.