

Household Saving Behavior and Social Security Privatization*

Alisdair McKay*

December 17, 2011

Abstract

I develop a general equilibrium model of saving behavior in which the quality of financial decisions is endogenously determined by the incentives to exert effort in learning about financial opportunities. The model generates predictions for asset holdings, asset market participation and returns across households. I apply the model to analyze a social security privatization reform. Privatization affects household search effort, asset market participation and the competitiveness of the asset market. Privatization reduces average welfare and this reduction is somewhat magnified by the search frictions. While some have suggested that household decision making could be important for the consequences of privatization, my analysis does not bear this out.

* I would like to thank Per Krusell, Nobu Kiyotaki, Ricardo Reis and Esteban Rossi-Hansberg for advising me in this work and seminar participants and others for many helpful comments and suggestions.

*Boston University. amckay@bu.edu

1 Introduction

Standard models of household consumption and savings assume that all households earn the marginal product of capital on their savings. This assumption is at odds with the evidence on the actual financial arrangements of households, which shows that a substantial number of households seem to have difficulty choosing a portfolio. Evidence of these difficulties takes several forms. For example, households may not allocate any savings to equities or they may hold under-diversified portfolios (Haliassos and Bertaut, 1995; Calvet et al., 2007). Surveys of financial literacy have also found that many households do not understand some fundamental financial concepts such as the difference between bonds and stocks (van Rooij et al., 2007). Other studies have found that those households that spend more effort planning for retirement reach retirement age with more wealth (Ameriks et al., 2003; Lusardi and Mitchell, 2007). In addition, researchers have found that experimental subjects have difficulty making sound financial decisions even when there is a clear normative ranking of the available choices (Choi et al., 2010). Many of these studies find that households with higher levels of income, wealth and education have more success in making sound financial decisions.

One way of understanding this empirical evidence is to view managing a portfolio as an activity that requires effort, with the incentive to devote effort varying across households. For example, households with high levels of wealth have more to gain in absolute terms from improving the return on their portfolios. Alternatively, highly educated households may be better able to assess the various risks and trade-offs that arise in choosing a portfolio. In this paper, I develop a general equilibrium model of household saving behavior in which households must exert effort to learn about the available investment opportunities by searching for high returns. In the model, a household can raise the expected return on its portfolio by devoting more effort to search. The benefit of search exists because there is dispersion in the rates of return offered by financial intermediaries. When households are imperfectly informed, intermediaries can still attract savings even if they are not offering the highest return, but intermediaries that offer higher returns will attract more savings. Therefore, intermediaries face a trade-off between the number of savers they will attract and their profit margin. In the model, these competing forces balance in such a way that intermediaries choose to offer a range of returns, which gives rise to an endogenous distribution of offered returns that depends on

the search and saving behavior of households.

The search friction is meant to capture a relationship between the effort households exert managing their finances and the return earned on their savings. Search frictions are a convenient way of modelling these relationships, but in reality there are more forces at work. For instance, some households have an imperfect understanding of the terms of the financial contracts or the principles of portfolio management. The point of view I take here is that households could learn these things if they devoted more effort.

To build the model, I draw on two literatures. The model is based on a heterogeneous-agent life-cycle savings model in the style of Bewley (undated), Huggett (1993) and Aiyagari (1994). It is natural to model household financial decisions within this framework because financial choices are non-linear functions of household assets, which means the distribution of financial outcomes will depend on the distribution of wealth. Given the relationship between financial outcomes and wealth, it is important to use a modeling framework that captures the heterogeneity in household assets. I modify the Bewley-Huggett-Aiyagari framework to include a search friction in the asset market. In the model, intermediaries post rates of return on risk-free assets and households choose how much time to spend searching among the offers for a high rate of return. To generate dispersion in returns, I use insights developed in the literature on search and equilibrium price dispersion, in particular from the work of Butters (1977) and Burdett and Judd (1983).¹ The other side of the asset market, in which intermediaries interact with production firms, is frictionless and does not play a major role in the analysis.

The model generates predictions for three aspects of household saving behavior that are absent from the standard model with a frictionless asset market. First, the model generates predictions for the amount of time that households spend managing their finances. I use data from the American Time Use Survey (ATUS) to calibrate the model and check these predictions. Second, the model generates predictions for asset market participation. In the model, non-participants are households that choose not to search or fail to

¹Drozdz and Nosal (2008) also embed a Burdett and Judd setting within a model of household saving behavior to study the market for unsecured borrowing. In their model, banks target offers to specific types of households and the number of offers that a household receives depends on how intensely banks are targeting those households and not on the household's search effort. Carlin and Manso (2011) develop a search theoretic model of the market for mutual funds in which expert consumer are perfectly informed and non-expert consumers choose at random. Their focus is on the incentives of the mutual fund industry to obfuscate the choice set and not on household saving behavior.

find an offer when they do search. I compare the model's predictions to data from the Survey of Consumer Finances (SCF). Third, the model generates predictions for the distribution of returns. These predictions do not have a clear empirical counterpart, but I use the distribution of fees on S&P 500 index mutual funds to calibrate and check the model. The model performs well on many of these dimensions.

The search friction also has implications for the distribution of wealth. Bewley-Huggett-Aiyagari models traditionally have had difficulty explaining the extreme skewness of the distribution of wealth. Cagetti and De Nardi (2006), Campanale (2007) and Benhabib et al. (2011) have shown that heterogeneity in household savings technologies can generate a more realistic distribution of wealth. While preceding work has relied on exogenous variation in the rates of return that households earn on their savings, the model presented here offers an endogenous mechanism that produces heterogeneity in returns. The mechanism at work was first pointed out by Arrow (1987): when households can pay to acquire information that will raise returns, wealthy households will acquire more information and earn higher returns, which leads to a more concentrated distribution of wealth. Indeed, the model predicts that wealthy and high-income households will be more likely to participate in the asset market and earn higher returns conditional on participation. This prediction is the result of a scale effect: as a household accumulates wealth, its incentive to search increases and it will earn higher returns on average. I am able to quantitatively explore the role of Arrow's mechanism in shaping the distribution of wealth. I find that the search friction does produce additional skewness in the distribution of wealth, raising the share of wealth held by the top quintile by 4.5 percentage points. I compare these predictions to those of a model with a fixed cost of asset market participation. That model does not generate as much wealth inequality because there is not heterogeneity in returns among those who participate in the market.

In an application of the model, I analyze the consequences of social security privatization in an environment in which households have difficulty allocating savings to the best investment opportunities. Many proposals for social security reform give individual households a larger role in managing their social security savings, but the empirical household finance literature raises questions about how well-prepared households are to take on this increased responsibility. To my knowledge, the literature on social security reform has

not formally analyzed how household financial decision making frictions affects the welfare implications of privatization.²

There are two views on how household financial decisions might affect the consequences of introducing private social security accounts. One perspective is that some households will make poor choices for their private social security accounts and will have few savings with which to retire.³ Another perspective is that fixed costs in portfolio management exacerbate the distortions caused by the social security system because they create an economy of scale in saving. If the system were privatized, the benefit of private saving would rise for two reasons. First, there is the usual reason that it is more important to save to provide for one's own retirement. Second, the accumulation of private savings would lead some households to overcome the fixed costs of managing their portfolios and realize a higher return on their savings.⁴

The model is able to generate both of the effects mentioned above. On the one hand, some households will have poor financial outcomes either due to a lack of effort or due to bad luck. On the other hand, privatization increases the size of the household's portfolio and strengthens the incentive to search, which raises their expected returns above what they would have earned if they did not adjust their search effort.

In the context of this model, social security reform can affect welfare through an additional channel: the competitiveness of the asset market. Privatization leads most households to increase their search effort as they hold more wealth. This increase in search effort has two partially offsetting effects on the distribution of offered returns. For those households who were already participating in the asset market, they are now more likely to have multiple offers and the higher degree of competition motivates firms to make more attractive offers. However, there are some households that were previously non-participants that enter the market after the reform. As these new participants are on the margin of non-participation, they choose relatively low levels of search effort, which motivates firms to make less attractive offers. The implications of privatization for the competitiveness of the asset market are therefore ambiguous.

²Some notable papers that analyze social security reform with standard assumptions about asset markets are Imrohorglu et al. (1995); Huang et al. (1997); Kotlikoff et al. (1998); Conesa and Krueger (1999); Huggett and Ventura (1999); Storesletten et al. (1999); Nishiyama and Smetters (2007), and Fuster et al. (2007). Hurst and Willen (2007) study reforms in an environment in which households earn different returns on savings than they pay to borrow.

³This view has been expressed by Diamond and Orszag (2004).

⁴A similar point has been made by Feldstein and Liebman (2002).

To explore these effects, I simulate a partial privatization of the social security system in which benefits and payroll taxes are reduced by roughly 50% and I compute a 150-year transition path from the initial steady state to the privatized steady state. I find that social security privatization leads households to devote more effort to search on average and leads firms to make more competitive offers, but the magnitude of the latter effect is small. The partial privatization produces an average welfare loss equivalent to 1.06% of consumption. To assess the impact of the search friction on this result, I also perform the same experiment in a benchmark economy that is identical except that households earn the marginal product of capital on their savings without having to exert any effort. The average welfare loss from privatization in the benchmark economy stands at 0.88% of consumption.⁵ In the long run, after the transition costs have been paid, privatization leads to welfare gains. Here the average welfare gain is smaller when the search friction is included, standing at 6.1% of consumption rather than 6.7% in the benchmark model. My interpretation of these findings is that the incorporation of frictions into household financial decision making does not fundamentally alter the analysis of social security privatization and so the analysis with the standard model is robust in this regard.

The next section introduces the model and Section 3 discusses its calibration. Section 4 describes the model's steady state and compares the model's prediction to the data. Section 5 presents the privatization experiment and section 6 concludes.

2 Model

This section presents the model environment, the household and firm decision problems and defines the equilibrium concepts.

⁵That the privatization experiment produces a welfare loss in the benchmark economy confirms the findings of Nishiyama and Smetters (2007). Similarly, Conesa and Krueger (1999) show that a majority of voters would reject social security reform that would phase out the system.

2.1 The environment

2.1.1 Population, preferences and endowments

There is a continuum of households that follow a life cycle. A household of age i survives to age $i + 1$ with probability ν_i with $\nu_T = 0$ for some terminal age T . Households younger than age $T^R < T$ are considered working-age and have positive labor productivity. Households of age T^R and older are considered retired, do not have productive labor endowments and draw retirement benefits from the social security system. When a household dies, it is replaced by N children that inherit its assets. The number of heirs increases with the household's age at death according to $N = (1 + \gamma)^i$, where i is the parent's age at death. Let Υ_t^i be the mass of households of age i at time t . The population-age structure evolves according to

$$\Upsilon_t^i = \begin{cases} \nu_{i-1} \Upsilon_{t-1}^{i-1} & \text{for } i > 1 \\ \sum_{j=1}^T (1 - \nu_j) (1 + \gamma)^j \Upsilon_{t-1}^j & \text{for } i = 1. \end{cases}$$

In this formulation, population growth depends on the distribution of household ages. Once the age distribution has settled down to its ergodic distribution, the population will grow at rate $(1 + \gamma)$ per period and this is the case considered throughout the paper.

Parents are altruistic and view their heirs as an extension of themselves. I use the following dynastic preference structure

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{[c_t^\chi (1 - n_t - s_t)^{1-\chi}]^{1-\rho}}{1 - \rho},$$

where c_t , n_t and s_t are consumption, labor supply and search effort in period t .⁶

Households receive stochastic endowments of labor productivity. A household's labor productivity depends on its age and on an idiosyncratic shock, which is denoted $\hat{\epsilon}$. $\hat{\epsilon}$ follows a Markov chain with m states.

The household's age and productivity shock make up the household's exogenous state, which is denoted by

⁶De Nardi (2004) has shown that intergenerational linkages and a voluntary bequest motive are important for generating realistic lifetime savings profiles and matching the upper tail of the wealth distribution. I have also considered a formulation of preferences in which discounted expected utility is weighted by the size of the dynasty at each date. That model produces similar results, but produces lifetime saving profiles with unrealistic saving rates for retired households.

$\epsilon = (i, \hat{\epsilon})$. The labor productivity of a household of type ϵ is then given by $y(\epsilon)$. The exogenous transitions can be written as a single Markov chain with Tm states and transition matrix Γ . Finally, the notation $N(\epsilon_{t+1}, \epsilon_t)$ is used to denote the change in the household's size from one period to the next. $N(\epsilon_{t+1}, \epsilon_t)$ is equal to one throughout a household's lifetime and equal to $(1 + \gamma)^i$ if the household dies at age i .

2.1.2 Technology

At each date there is a continuum of firms operating a Cobb-Douglas production technology that combines capital and labor to produce a composite good according to $k^\alpha \ell^{1-\alpha} + (1 - \delta)k$. Let μ_t be the mass of firms at date t normalized by the mass of households at date t . Importantly, households do not have direct access to the production technology and must invest through firms. In addition to the production technology there is a storage technology that yields a net return of zero and is not subject to depreciation. This storage technology will provide a reservation return in the asset market. There is no aggregate risk in the economy.

2.1.3 Market structure

There are markets for labor and capital services in which firms rent capital and labor from households. Households are unable to borrow. The market for labor is Walrasian and is cleared by the wage w_t at each date t . Importantly, firms choose their labor inputs after capital is in place so labor mobility equalizes the capital-labor ratio across firms and therefore the marginal product of capital is common across firms. The marginal product of capital will be denoted by A_t .

In reality, households rarely interact with the ultimate users of their savings, but instead these relationships are intermediated by one or more financial institutions. Suppose the market in which the financial intermediary lends to production firms is frictionless, then the return on funds in this market will be equal to the marginal product of capital for production firms. The intermediaries will then interact with households as if they were directly investing in capital and earning the marginal product of capital. Therefore, the model abstracts from this second market and proceeds as if households interacted directly with production firms.

The market for capital services is where the model differs from a typical Bewley-Huggett-Aiyagari model. There is a search friction such that households cannot observe all available capital rental rates. Firms commit to and post returns and households search among those returns. There is no uncertainty in production and all offered returns are risk-free. Therefore, a household that encounters multiple firms during search will select the highest offered return and invest all of its savings with that firm.

A household chooses an amount of search effort, s , which generates a stochastic number of offers, j , with probability $q(j; s)$. A household with assets a and exogenous state ϵ will search an amount $s(a, \epsilon)$ and so the compact notation $q(j; a, \epsilon) \equiv q(j; s(a, \epsilon))$ will be used. It is assumed that the sequence of q 's satisfies the following assumption, which guarantees that the firm's problem is well-defined.

Assumption 1. *For any s , $\sum_{j=1}^{\infty} jq(j; s)$ is finite.*

When a household meets a firm it receives one random draw from the distribution of offered returns, which has cumulative density function $F(r)$. Firms must pay a fixed cost ψ to post a return. Finally, matches last for a single 5-year period after which the relationship is dissolved.

2.1.4 Government and social insurance

There is a government that levies taxes, consumes resources and distributes social security benefits. Following Castaneda et al. (2003), a household's social security benefit depends on its final working-age labor productivity. Therefore, the benefit can be written as $B_t(\epsilon)$ where it is assumed that $B_t(\epsilon) = 0$ for working-age households and that the idiosyncratic component of ϵ does not change during retirement. The dependence of the benefit on final working-age productivity is able to capture some of the progressive nature of the social security system without introducing an additional state variable.

Social security benefits are funded through a dedicated payroll tax as is the case in the US. The tax rate is denoted τ_t^y . In addition, an income tax, τ_t , is used to fund an exogenous sequence of general government expenditures $\{Q_t\}_{t=0}^{\infty}$. These expenditures have no role other than absorbing tax revenues. The government has separate budgets for general expenditures and transfer payments and each budget is balanced period-by-period. I also allow for a consumption tax, τ_t^c , which will be used in the privatization experiment, but is

set to zero in the steady state.

2.1.5 Discussion

There are several aspects of the model environment that deserve comment. First of all, the assumption that all assets are risk-free is an important simplification. These risk-free assets are perhaps best viewed as the certainty-equivalent values of more complicated portfolios that an intermediary may offer. As such, the model is able to capture the heterogeneity in expected returns across households, but it cannot capture differences in diversification.

Second, one might ask why households do not delegate their search to a well informed agent? Even if households hire financial advisors to handle their portfolios, some degree of information acquisition effort is inescapable as the household must choose an advisor. There is always a first link between the household and the financial system and this is the link that is modeled here.

Third, the search friction is a source of increasing returns to wealth and, as a result, the household's value function may not be concave. I have experimented with introducing actuarially-fair lotteries that allow households to smooth out these non-concavities. In the computed equilibrium of that model, only 3% of households trade lotteries and the results are nearly identical to those when the lotteries are omitted. Therefore, I continue without lotteries for the sake of simplicity.

Finally, I assume that matches only last for a single model period, which I calibrate to be five years. While the assumption that matches last for a single period is a simplification, modeling financial decision making as a once-in-a-lifetime event would also be unsatisfactory as households devote effort to financial decisions throughout their lives.⁷ The truth surely lies somewhere in between with households engaging in some amount of learning in early years after which they must continually respond to changing circumstances and opportunities.

⁷For example, Figure 3 suggests that households actually devote more and more time to financial decisions as they age.

2.2 Decision problems

2.2.1 The household's problem

The household must choose consumption, savings, search effort and labor supply to maximize expected utility subject to a budget constraint, a borrowing constraint and the distributions of returns and exogenous states. Households are uncertain about the return they will earn because they do not know how many firms they will meet nor what returns those firms will offer. A choice of search effort s , therefore, generates a distribution over returns that has cumulative distribution function

$$G_t(r; s) \equiv \sum_{j=0}^{\infty} q(j; s) [F_t(r)]^j. \quad (1)$$

Using this notation the household's problem can be written recursively as

$$V_t(a, \epsilon) = \max_{c, a^+, s, n} \{u(c, 1 - n - s) + \beta E_t [V_{t+1}(a', \epsilon')]\}$$

such that

$$(1 + \tau_t^c) c = a + (1 - \tau_t - \tau_t^y) n y(\epsilon) w_t + B_t(\epsilon) - a^+$$

$$a' = [1 + r' (1 - \tau_{t+1})] \times a^+ / N(\epsilon', \epsilon)$$

$$r' \sim G_t(\cdot; s)$$

$$\epsilon' \sim \Gamma(\cdot; \epsilon)$$

$$0 \leq s, \quad 0 \leq n, \quad a^+ \geq 0.$$

The household sets aside savings of a^+ , which grow to $a' = [1 + r' (1 - \tau_{t+1})] \times a^+$ by next period if the household does not die. If the household dies, the bequest is divided among the $N(\epsilon', \epsilon)$ heirs. The return r' is drawn from the distribution with CDF G that depends on the household's search effort and the distribution of offered returns. Let the household's decision rules be denoted by $h_t(a, \epsilon) = a^{+*}$, $c_t(a, \epsilon) = c^*$, $s_t(a, \epsilon) = s^*$

and $n_t(a, \epsilon) = n^*$.

2.2.2 The firm's problem

The firm's problem is to choose a return to post that maximizes profits. Expected profits are given by⁸

$$\pi(r) = \left(1 - \frac{1+r}{A}\right) \sum_{\epsilon} \int h(a, \epsilon) \left\{ \mu^{-1} \sum_{j=0}^{\infty} j q(j; a, \epsilon) [F(r)]^{j-1} \right\} \Phi(da, \epsilon). \quad (2)$$

This expression is the product of the profit margin, $(1 - \frac{1+r}{A})$, and the expected assets that will be attracted when return r is posted. The latter involves multiplying the amount a household of type (a, ϵ) saves against the probability that the household will choose to do business with a firm posting return r , which depends on its level of search effort through the q 's. The term $h(a, \epsilon)$ is the amount of savings rented from a household of that type. The term in braces is the probability that the firm will meet and rent capital from a household of type (a, ϵ) when it posts return r , which takes the same form as in Burdett and Judd (1983). Finally, the integral is over the distribution of households over the state space, Φ .

2.3 Equilibrium

For a given joint distribution of household savings and search effort choices, the firms play a return-posting game. I begin by defining an equilibrium of the return-posting game before turning to the broader equilibrium of the model.

2.3.1 Firm equilibrium

Following Burdett and Judd (1983), I define an equilibrium of the return-posting game.

Definition For a given marginal product of capital, A , distribution of households over the state space, $\Phi(a, \epsilon)$, and decision rules, $s(a, \epsilon)$ and $h(a, \epsilon)$, a *firm equilibrium* is an offer distribution $F(\cdot)$ and a scalar π^* such that $\pi(r) = \pi^*$ for all r in the support of $F(\cdot)$ and $\pi(r) \leq \pi^*$ for all r outside the support of $F(\cdot)$.

⁸Time subscripts have been suppressed for clarity as the firm's problem is static although the reader should keep in mind that A refers to the marginal product of capital that will prevail in the following period.

Two important properties of any firm equilibrium are that a) the offer distribution $F(r)$ is continuous and b) $r = 0$ is in the support of $F(\cdot)$ with $F(0) = 0$. Burdett and Judd (1983) prove these properties for the case of homogeneous consumers. These basic properties of the firm equilibrium are unchanged by the introduction of heterogeneous households as long as the following assumptions are satisfied.

Assumption 2. *i . Some resources are invested after observing exactly one return:*

$$\sum_{\epsilon} \int h(a, \epsilon) q(1; a, \epsilon) \Phi(da, \epsilon) > 0.$$

ii . Some resources are invested after observing two or more returns:

$$\sum_{\epsilon} \int h(a, \epsilon) \sum_{j=2}^{\infty} q(j; a, \epsilon) \Phi(da, \epsilon) > 0.$$

iii . Total resources invested are finite:

$$\sum_{\epsilon} \int h(a, \epsilon) \Phi(da, \epsilon) < \infty.$$

When I turn to numerical solutions of the model I will specify a sequence $\{q(j; s, \epsilon)\}_{j=0}^{\infty}$ for which the first two elements of Assumption 2 hold as long as a positive mass of households chooses positive search effort. The implication is that the economy is not in an equilibrium in which firms post the reservation return and households do not search.

Proposition 1. *If Φ , $h(a, \epsilon)$ and $s(a, \epsilon)$ are such that Assumption 2 holds, then if $F(\cdot)$ is part of a firm equilibrium,*

1. $F(\cdot)$ is continuous
2. the support of $F(\cdot)$ starts at zero (reservation return),
3. the support of $F(\cdot)$ ends at some $\bar{r} < A - 1$,
4. $F(\cdot)$ is strictly increasing on $[0, \bar{r}]$.

Proof. See Appendix A. ■

2.3.2 Heterogeneous firms

The firm equilibrium relies heavily on the homogeneity of firms and their exact indifference across the support of the offer distribution. As it is unlikely that real-world firms or financial intermediaries are

completely identical, it is important to check that a small amount of heterogeneity among firms does not produce a substantially different offer distribution. Fortunately, the equilibrium is robust to small amounts of heterogeneity. Appendix B shows that if firm marginal productivities are distributed between \underline{A} and \bar{A} , then as \underline{A} and \bar{A} converge to A the offer distribution converges point-wise to the offer distribution that arises when firms are homogeneous with productivity A .⁹

2.3.3 Recursive equilibrium

I define an equilibrium of this economy recursively. Beyond the firm equilibrium, the rest of the recursive equilibrium definition is an extension of the same concept for a Bewley-Huggett-Aiyagari economy.

Definition For a given initial capital stock and a given initial distribution of households over the state space, a recursive equilibrium is a sequence of objects $\{ V_t, h_t, c_t, s_t, n_t, \Phi_{t+1}, F_t, \pi_t^*, \mu_t, w_t, A_t, K_{t+1}, L_t, \tau_t, \tau_t^y, \tau_t^c \}_{t=0}^{\infty}$ such that for each date t

1. given A_{t+1}, h_t, s_t , and $\Phi_t, \{F_t(\cdot), \pi_t^*\}$ is a firm equilibrium
2. firms have no incentive to enter or exit: $\pi_t^* = \psi$
3. V_t solves the consumer's problem with policy rules h_t, c_t, n_t and s_t
4. $G_t(\cdot)$ is generated from $F_t(\cdot)$ according to equation (1).
5. the distribution of households over (a, ϵ) -space evolves according to¹⁰

$$\Phi_{t+1}(\mathcal{A}, \epsilon') = \sum_{\epsilon} \Gamma_{\epsilon, \epsilon'} \int \int_{(1+r'(1-\tau_{t+1})) \times h_t(a, \epsilon) \in \mathcal{A}} G_t(dr'; s(a, \epsilon)) \Phi_t(da, \epsilon) \quad (3)$$

$$6. K_{t+1} = \sum_{\epsilon} \int h_t(a, \epsilon) [1 - q(0; a, \epsilon)] \Phi_t(da, \epsilon) - \psi \mu_t \sum_i \Upsilon_t^i$$

$$7. L_t = \sum_{\epsilon} \int y(\epsilon) n_t(a, \epsilon) \Phi_t(da, \epsilon)$$

⁹Bontemps et al. (1997) show a similar result in the context of a model of on the job search.

¹⁰The distribution with CDF G places a mass of probability on $r = 0$. The notation $\int f(r)G(dr)$ for some function $f(r)$, should be understood as $G(0)f(0) + \int_0^{\bar{r}} f(r)g(r)dr$, where $g(r)$ is the density of $G(r)$ on $(0, \bar{r})$.

$$8. w_t = (1 - \alpha) \left(\frac{K_t}{L_t} \right)^\alpha$$

$$9. A_t = \alpha \left(\frac{K_t}{L_t} \right)^{\alpha-1} + 1 - \delta$$

10. the government budget for general expenditure is in balance:

$$\tau_t \sum_{\epsilon} \int w_t y(\epsilon) n(a, \epsilon) \Phi_t(da, \epsilon) + \tau_t \sum_{\epsilon} \int \int r' h_{t-1}(a, \epsilon) G_{t-1}(dr'; s(a, \epsilon)) \Phi_{t-1}(da, \epsilon) = Q_t$$

11. and the social security budget is in balance

$$\sum_{\epsilon} \int \tau_t^y w y(\epsilon) n(a, \epsilon) + \tau_t^c c_t(a, \epsilon) - B(\epsilon) \Phi_t(da, \epsilon) = 0.$$

The aggregation of household savings into the capital stock, on line 6 of the definition, differs from the usual expression in two ways. First, there is an adjustment for households that do not meet a firm and use the storage technology. The fraction of households of type (a, ϵ) that use the storage technology is given by $q(0; a, \epsilon)$. Second, the fixed costs of return posting are paid out of household savings. There is a mass $\sum_i \Upsilon_t^i$ of households at date t , μ_t firms per household, and each firm has a fixed cost of ψ so the capital stock is smaller than the aggregate of household savings by an additional amount equal to the product of these terms.

The definition of a steady state is also a modified version of the definition for a Bewley-Huggett-Aiyagari economy. In particular, in a steady state the distribution over the state space, Φ , and the household decision rules are such that the same distribution Φ is generated in every period.

3 Calibration

I must specify a functional form for the search technology before discussing parameter values. I assume that offers arrive according to a non-homogeneous Poisson process during the time a household spends searching

so that

$$q(j; s) = \frac{(\theta(s))^j e^{-\theta(s)}}{j!}, \quad (4)$$

where $\theta(s)$ is the integral of the arrival rate from zero to s , which I assume takes the form

$$\theta(s) = \theta_1 \times \log(1 + \theta_2 \times s).$$

The firm's profit function from equation (2) then reduces to

$$\pi(r) = \left(1 - \frac{1+r}{A}\right) \sum_{\epsilon} \int h(a, \epsilon) \mu^{-1} \theta(s) \exp\{\theta(s) [F(r) - 1]\} \Phi(da, \epsilon)$$

and the distribution of returns that a household earns defined by equation (1) reduces to

$$G(r; s) = \exp\{\theta(s) [1 - F(r)]\}. \quad (5)$$

A fraction $\exp\{\theta(s)\}$ of households that choose a particular level of search effort, s , will fail to receive even one offer and will therefore not participate in the asset market.

Notice that $\theta''(s)$ is negative so there are decreasing returns to search effort in the sense that it becomes increasingly difficult to obtain new offers as one searches more. The two parameters θ_1 and θ_2 can be used to control both the average and marginal values of search.¹¹ I can choose how much time the agents devote to search and how many offers they will have on average given that amount of search, which will determine the level of competition among firms and therefore the distribution of offered returns. In the next sub-section, I use these predictions to calibrate these two parameters.

The model period is calibrated to be five years. Households are born at age 21 and survival probabilities are taken from the Social Security Administration (2007).¹² Households die with certainty at age 110 and

¹¹Notice that $\theta'(s) = \theta_1 \theta_2 / (1 + \theta_2 s)$ so the parameters θ_1 and θ_2 allow one to control the initial efficiency of search, $\theta'(0) = \theta_1 \theta_2$, by adjusting θ_1 and the degree to which this efficiency declines as one searches more by adjusting θ_2 .

¹²The Social Security Administration reports annual survival probabilities. To convert to a five year model period, I take the

are considered to be retired after age 65. Population growth is set to 1.27% per year, which matches the growth rate of the United States population from 1940 to 2000.

3.1 Calibration of endowments

For working-age households, log labor productivity is given by the sum of a life-cycle effect and an idiosyncratic, persistent shock. The life-cycle component is calibrated using labor income per hour from the Panel Study of Income Dynamics (PSID) for the nine age groups 21 - 25 through 61 - 65.¹³ The persistent shock follows a discrete approximation to an AR(1) process with autoregressive parameter 0.91 and normal innovations with mean zero and standard deviation 0.29. These parameters are taken from estimates by Heathcote et al. (2010).¹⁴ The process is discretized to seven points using the method of Tauchen (1986).

The labor productivity levels of parents and children are correlated and this correlation has implications for the degree of wealth inequality in the economy. When a model household is born, it inherits its parent's final labor productivity with probability 0.87 and otherwise receives a new draw from the ergodic distribution. The value 0.87 is chosen as this generates a correlation of 0.6 between the average lifetime productivity levels of parents and children, which is guided by estimates of a correlation of around 0.6 between the lifetime earnings of fathers and sons (Mazumder, 2005; Gouskova et al., 2010).

3.2 Calibration of preference and production parameters

The discount rate is set to 0.859 in order to match the capital-output ratio. The coefficient of relative risk aversion is set to 2. A value of 0.373 is chosen for χ to match an average labor supply of 0.35, which is the fraction of time devoted to work in the ATUS data. A value of 0.36 is chosen for α to match the capital share. Finally, the depreciation rate is set to 0.294 in order to match an annual investment rate of 8%.

product of the five annual values.

¹³See Appendix D for details.

¹⁴The estimates are transformed in two ways. First, Heathcote et al. allow the variance of the innovations to change over time and I use the average variance over their 1969 - 2000 sample period, which produces a value of 0.0139. Their estimate of the autoregressive coefficient is 0.9733. To convert these values to a five-year model period, I simulate annual data, take five year averages and estimate the parameters reported above using ordinary least squares.

3.3 Calibration of the search parameters

The search efficiency parameters, θ_1 and θ_2 , are calibrated to two moments. First, the average fraction of time that households devote to search matches the average fraction of time devoted to managing household finances in the ATUS.¹⁵ This figure is three minutes per day or 0.3% of time not devoted to personal care. Second, I define an asset management fee as the difference between the marginal product of capital and the offered return. I calibrate the model so that the median fee matches the median fee in a sample of 109 S&P 500 index mutual funds, which is 64 basis points per year.¹⁶ S&P 500 index funds are by no means the universe of investment opportunities, but are one way of comparing the return dispersion in the model to the return dispersion in the data. Despite holding very similar portfolios, S&P 500 index funds do not all charge similar fees and Hortacsu and Syverson (2004) have argued that search frictions are important to sustaining the price dispersion in this market.

The calibration strategy does not explicitly target the dispersion in offered returns, but the degree of dispersion in returns will be determined by the degree of competition in the market, which also determines the median fee.

The fixed costs of return posting, ψ , can be normalized to unity because any change in fixed costs is exactly offset by a change in the mass of firms μ . Given the other parameters of the model, there is a certain mass of total profits and changing the per-firm fixed cost only has implications for the mass of firms that enter the market.

3.4 Calibration of the government and social insurance parameters

The income tax is calibrated to the return-weighted average marginal income tax rate reported in Stephenson (1998) averaged over the years 1980-1994, which is 19.4%. Government expenditures, Q , are set to the level

¹⁵The ATUS asks respondents how much time during the reference day they spent on “household financial management” and “banking and using financial services.” I include both of these categories in my calculations.

¹⁶Data are from CRSP for retail mutual funds in 2005 and only retail funds are included in the sample. The cost of holding a mutual fund is a combination of the annual expenses and loads. The fees are calculated as annual expenses plus one seventh of loads, which annualizes the loads over a hypothetical seven year holding period.

symbol	description	value	target	target value	model value
β	discount factor	0.859	capital-output ratio	3.32	3.32
ρ	risk aversion	2	–	–	–
χ	labor supply	0.373	average hours	0.35	0.35
α	capital share	0.36	capital share	0.36	0.36
δ	depreciation rate	0.294	investment rate	0.08	0.08
γ	population growth	6.5%	–	–	–
θ_1	search efficiency	1.25	median fee	64 basis points	68 b. p.
θ_2	search curvature	2600	avg. search effort	3.0×10^{-3}	3.0×10^{-3}
ψ	fixed cost	1	normalization	–	–
τ	income tax rate	0.19	avg. marginal tax rate	–	–
τ_c	consumption tax rate	0.0	–	–	–
τ_y	payroll tax	0.093	soc. sec. budg.	–	–
Q	gov. cons. (Q/Y)	0.19	balanced budget	–	–

Table 1: Baseline calibration for a five year model period.

that absorbs these tax revenues, which is 19.4% of output.¹⁷

The social security system is calibrated to roughly match the system currently in place in the US. The US system makes payments conditional on average earnings over the household’s lifetime, but the system is not directly implemented in the model because that would require an additional continuous state variable in the household’s problem, which would be computationally costly. Instead, benefits depend on final labor productivity. Benefit levels are chosen so that applying the US benefit formula to simulated earnings histories generates the benefit levels used to solve the model. The replacement rates that result from this procedure range from 78% for low-productivity types to 13% for high-productivity types. Finally, the social security tax rate is set to balance the social security program budget. The level of social security tax needed is 9.3%.

3.5 A benchmark calibration

As a point of reference, I also compute solutions for a benchmark economy without the search friction. As $\theta_1 \times \theta_2$ goes to infinity, the search friction disappears and the model becomes a more standard Bewley-Huggett-Aiyagari model in which all households earn the marginal product of capital on their savings. This model is calibrated to match the same moments as the full model.¹⁸

¹⁷This is not a coincidence, but follows from $Q = \tau Y$. In the data, government consumption and gross investment were, on average, 20.3% of gross domestic product between 1960 and 2009.

¹⁸In the model without the search friction, β is set to 0.843, χ is set to 0.376 and τ_y is set to 0.092. The social security benefits are re-calibrated using the procedure described above. All other parameters are left unchanged.

3.6 Computation

The algorithm to solve the model is based on the same logic as the usual one for a Bewley-Huggett-Aiyagari economy. Appendix C provides details.

4 Steady-state results

I now describe the stationary equilibrium of the model. I group the model's predictions into four categories: the choice of search effort, asset market participation, the dispersion of returns, and the distribution of wealth. In each category, I compare the model's predictions to the available data.

4.1 The choice of search effort

The first panel of Figure 1 shows the function $\theta(s)$, which is also the expected number of offers that the household receives. The second panel of the figure shows the expected return that the household earns as a function of its search effort. There are decreasing returns to search effort for two reasons, first, as shown in Panel A., there are decreasing returns to search effort in terms of generating new offers, and second, there are decreasing returns to new offers because an additional offer is only valuable if it exceeds all others the household has already encountered.

Figure 2 shows the search decision rules for low-productivity and high-productivity households in the 41 - 45 age group. For high asset levels, low-productivity households search more than the high-productivity households because their opportunity cost of search is lower. For low asset levels, the low-productivity households choose not to search. High-productivity households, however, exert positive search effort at all asset levels because a high-productivity household that enters the period with no assets will still save enough out of that period's labor income to make it worthwhile to search. It is clear that a household's search behavior is increasing in initial assets. On average, high-productivity households accumulate more assets than low-productivity households and search more as a result. Higher levels of search effort lead to higher levels of asset market participation and higher average returns conditional on participation.

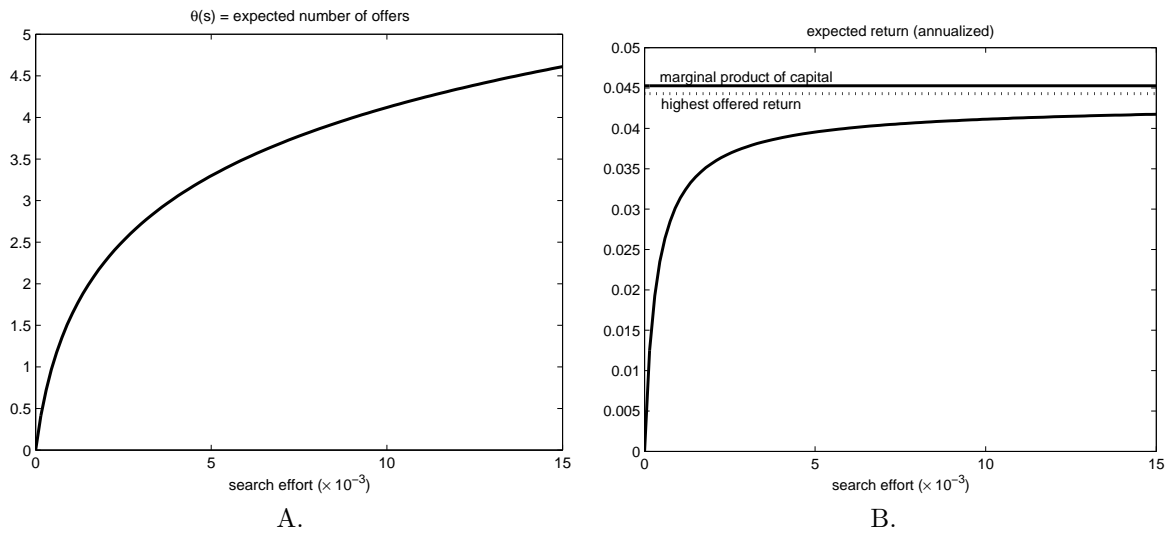


Figure 1: Returns to search. Panel A shows the expected number of offers a household receives for a given amount of search effort. Panel B show the expected return that a household earns after a given amount of search effort.

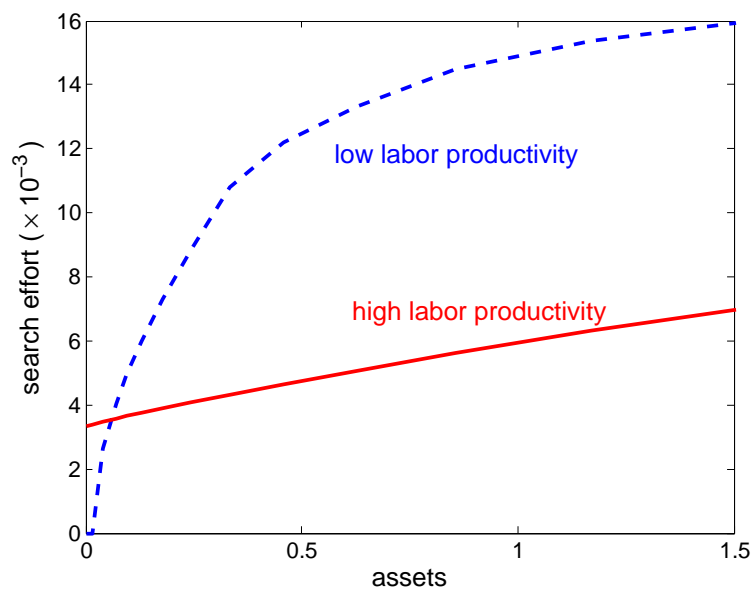


Figure 2: Search decision rules for lowest and highest productivity households in age group 41 to 45.

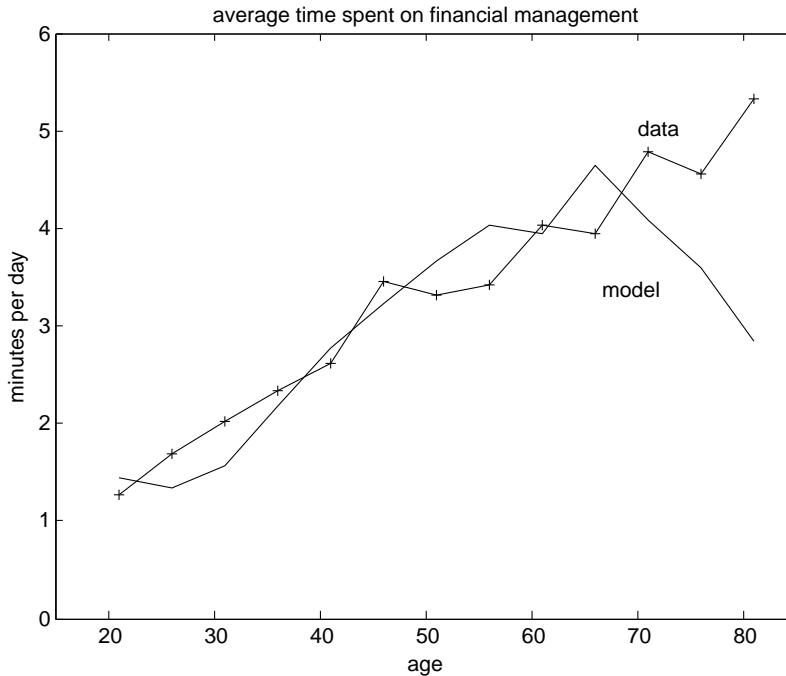


Figure 3: Simulated and empirical life-cycle profiles of search effort. Time use data are from the American Time Use Survey for years 2003 - 2006.

Figure 3 shows average search effort over the life cycle. Households devote more time to search as they accumulate higher levels of savings in anticipation of retirement despite the countervailing force of the increase in their labor productivity and the opportunity cost of search. Search effort is highest in the first period of retirement because households still have large asset positions and the opportunity cost of time has fallen. Households begin to search less as they run down their assets in retirement. I compare this life-cycle profile to the ATUS data.¹⁹ The model behaves well for working-age households and those in early retirement, but the decline in search effort among retirees runs counter to the evidence. The model only captures a subset of the activities involved in personal financial management and abstracts from estate planning and portfolio reallocation, which may become more important as a household grows older.

¹⁹The model is calibrated to match the average search effort across all households so the model fits that aspect of the data by construction.

4.2 Asset market participation

In the model, households may not participate in the asset market for two reasons. First, households may devote no effort to search, in which case they have no chance of encountering a firm. Second, they may devote effort but fail to meet a firm in the stochastic search process. Both of these effects together produce a steady-state asset market participation rate of 60%. This non-participation behavior is reminiscent of the limited stock-market participation observed by Mankiw and Zeldes (1991), Haliassos and Bertaut (1995), Vissing-Jorgensen (2002) and others. According to the 2004 Survey of Consumer Finances (SCF), 35% of the population holds no financial assets other than checking and savings accounts. If we include CDs and savings bonds, this figure rises to 42%. That means that only 58% of households own any of a 401(K), IRA, mutual fund, bond, stock, trust, or annuity. Moreover, only 50% of households own equities either directly or indirectly. Thus, the model's prediction of a 60% asset market participation rate is quite plausible. One feature of the stock market participation puzzle is that there are households with considerable savings that do not hold stocks. The model does predict that there will be non-participants even among the wealthiest households because those households may search, but fail to find a firm. Quantitatively, however, the model predicts very high rates of participation among the wealthy and most of the non-participants have little or no wealth.

Figure 4 shows the life-cycle profile of asset market participation, which reflects the profile of search effort. The participation rate rises sharply in middle age and declines somewhat in retirement. The figure also presents the empirical profile for a broad definition of asset market participation that classifies a household as a participant if it has any account or asset other than a checking or savings account.²⁰ Like the simulated life-cycle profile, the empirical profile peaks in middle age, but the increase is more gradual. Moreover, the decline in financial participation among retired households occurs at younger ages than the model predicts. These data should be interpreted with some care as some differences across ages may reflect cohort rather than age effects. Within an age group, the model predicts that wealthier households and households with high incomes are more likely to participate. These results are in line with the empirical findings of Vissing-

²⁰Specifically, participants have at least one of 401(K), IRA, mutual fund, bond, stock, trust, or annuity.

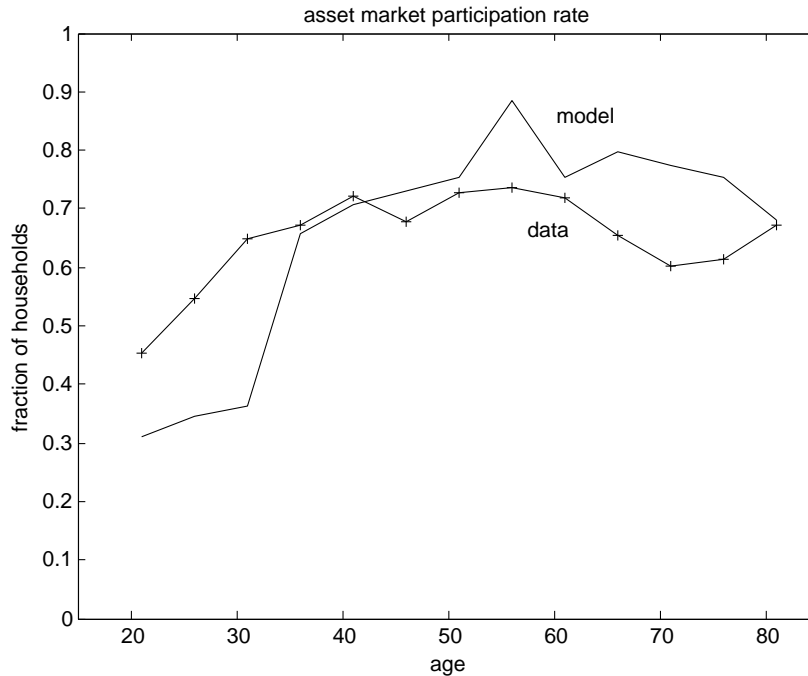


Figure 4: Simulated and empirical life-cycle profiles of asset market participation. Participation is defined as having any account or asset other than a checking or savings account.

Jorgensen (2002) and Calvet et al. (2007).

4.3 The dispersion of returns

I now turn to the dispersion of returns. I begin with the distribution of returns across firms (i.e. the offer distribution) and then look at the distribution of returns across households.

Panel A.i. of Figure 5 shows the model's distribution of offered returns. The distribution shows the usual properties of price distributions from a Burdett and Judd style search model: The support of the distribution extends all the way down to the reservation option, the density is concave, and the highest offered return is below the marginal product of capital. Despite the long tail of bad offers, the market is quite competitive with a large fraction of offers close to the marginal product of capital.

As a check on the model, Panel A.ii. of Figure 5 plots the distribution of fees on a sample of 109 S&P 500 index mutual funds. The lowest fee in the data is nine basis points compared to ten in the model.²¹ In

²¹The support of the offer distribution in the model ends ten basis points below the marginal product of capital.

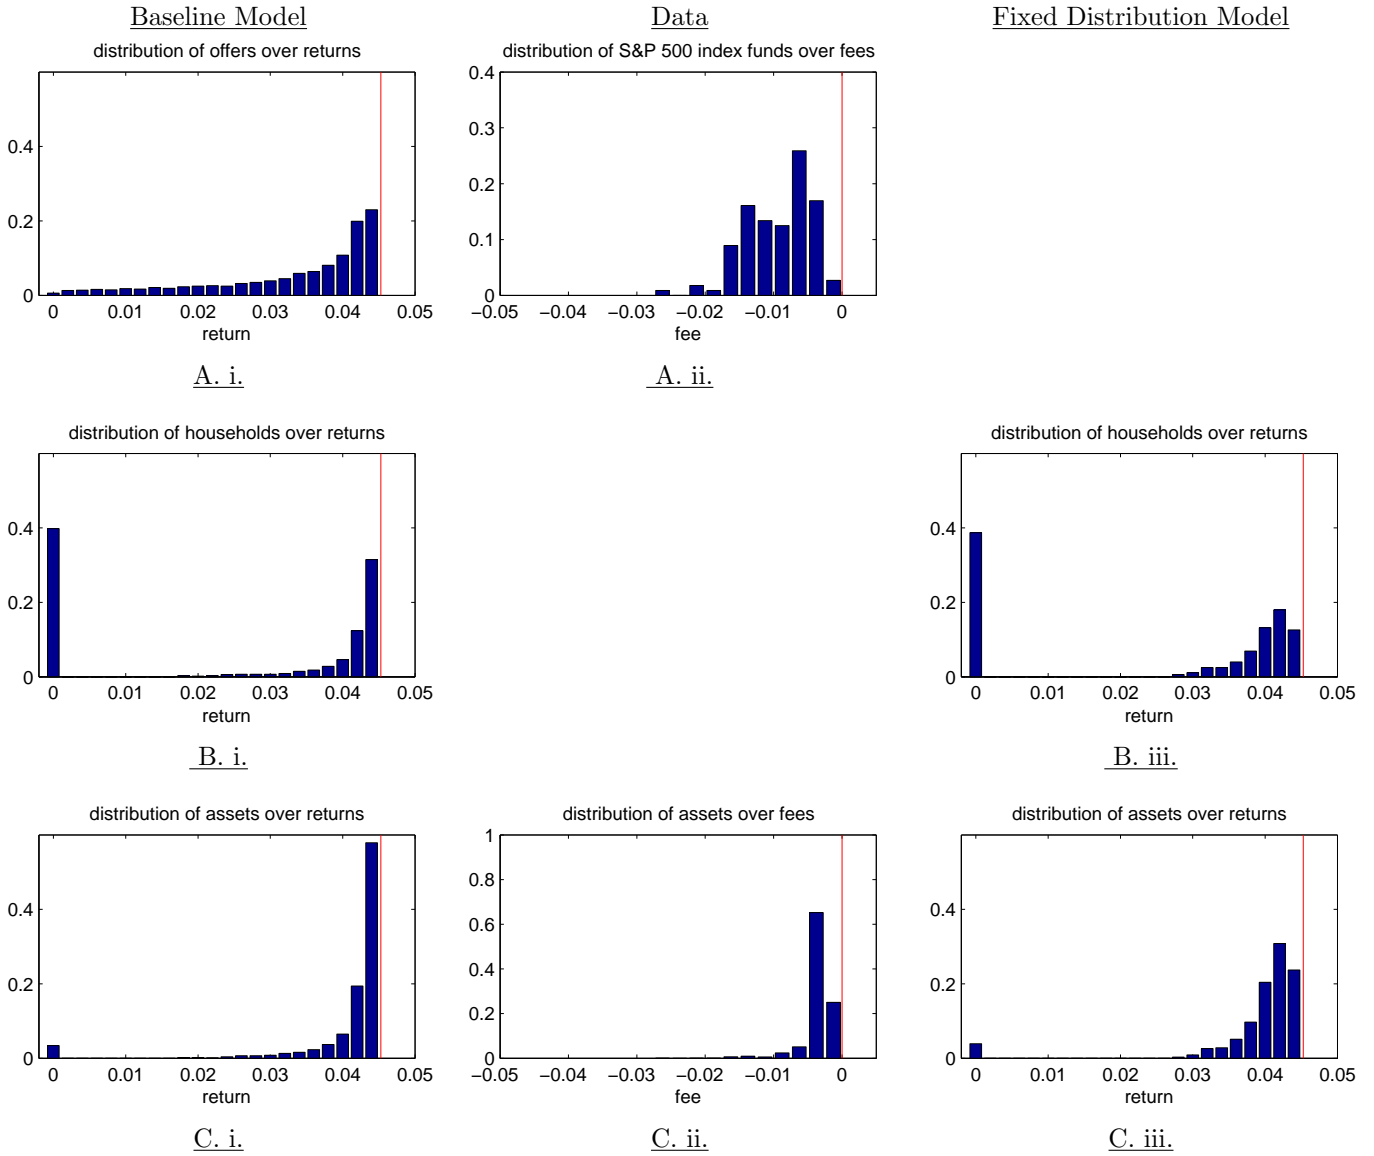


Figure 5: Figures on the left show the distribution of firms, households and assets over rates of return from the baseline model. The thin vertical line represents the marginal product of capital. The returns shown are in annual terms. The center column of figures shows the distribution of S&P 500 mutual funds and assets over fees. The horizontal axis has been reversed so that better offers are on the right. Figures on the right are from the model with the offer distribution derived from the distribution of fees on S&P 500 index funds (see section 4.5).

both the data and the model, the bulk of the offers have fees less than 150 basis points.

Panel B.i. of Figure 5 shows the distribution of households over returns. Conditional on participation, households are more concentrated near the marginal product of capital than the offers are. The marginal product of capital is 4.53% and the mean return among participants is 4.08%.²² However, the distribution of households over returns is heavily skewed and the median return is 4.34%.

Panel C.i. of Figure 5 shows the distribution of assets over returns, which is even more concentrated than the distribution of households because wealthier households tend to search more. It is also clear from the distribution of assets that non-participants hold little wealth. Panel C.ii. shows that the distribution of assets over the S&P 500 mutual funds is similar in that it is highly concentrated on the good offers.

4.4 The distribution of wealth

As wealthy households tend to earn higher returns, the search friction generates additional skewness in the distribution of wealth. Table 2 shows the concentration of wealth in the data and the model. For the sake of comparison, the table also presents results for the same economy without the search friction. The search friction results in a more concentrated distribution of wealth, and the top quintile's share rises 4.5 percentage points and the wealth Gini rises to 0.78 from 0.74. Models of this type have difficulty matching the asset holdings of the top 1% of the distribution and the search friction only helps marginally in this regard.

As another comparison, Table 2 presents the distribution of wealth from a model with a fixed cost of participating in the asset market. In this alternative benchmark, a household earns the marginal product of capital if it pays the participation cost. The fixed cost is denominated in units of time and is calibrated to produce a 60% asset market participation rate. This model produces more inequality in wealth than the model without search frictions, but not as much as that with the search friction. In particular, the fixed-cost model does not produce as much dispersion in the right tail of the distribution as evidenced by the fact that the share held by the top quintile rises by 2.7 percentage points rather than 4.5 for the model with search frictions. This makes sense as in the fixed-cost model there is no heterogeneity in financial outcomes among

²²All returns are in annual terms.

	Gini	Top Groups		Quintiles				
		1%	5%	1st	2nd	3rd	4th	5th
Data	0.80	34.7	57.8	-0.3	1.3	5.0	12.2	81.7
With Search Friction	0.78	16.5	45.0	0.0	0.1	3.7	14.3	81.9
Without Search Friction	0.74	14.6	40.6	0.0	0.1	5.6	16.1	77.4
Fixed Participation Cost	0.77	15.1	42.2	0.0	0.0	4.2	15.7	80.1
Exog. Offer Distribution	0.78	16.5	45.1	0.0	0.1	3.7	14.2	82.0

Table 2: The distribution of wealth. Data calculations are by Budria Rodriguez et al. (2002) using the 1998 SCF.

participants and households in the right tail are very likely to participate.

4.5 Return dispersion and saving behavior

Comparing Panels A.i. and A.ii. of Figure 5, one can see that the model does not exactly match the distribution of offered returns. As a robustness check, this section explores how the differences in the offer distribution affect household saving behavior by solving a version of the model in which the offer distribution is taken to exogenously reflect the distribution of fees on S&P 500 index funds. I smooth the observed data on fees with a kernel density estimate to generate an offer distribution from the data. I recalibrate the model to match the moments reported in section 3 with the exception that θ_1 is left at its original value.

The results of this exercise reveal that the shape of the offer distribution has only a small impact on household saving behavior. For example, Figure 6 plots the life-cycle profiles of savings, search effort and asset market participation when the offer distribution is fixed exogenously along with the same profiles from the baseline model with an endogenous offer distribution. The first panel of the figure also includes the prediction of the benchmark model with a frictionless asset market. In this panel, one can see that the two models with frictional asset markets produce extremely similar predictions for the profile of median assets as the two lines are indistinguishable in the figure. In all three panels, the exact shape of the offer distribution appears to make little difference to the way households behave. Moreover, the last row of Table 2 shows that the distribution of wealth is little changed when the offer distribution is exogenous. The principal difference between the two offer distributions is that the endogenous distribution has a long tail of bad offers. It turns out that this tail is not very important because most low returns will be rejected during the search process.

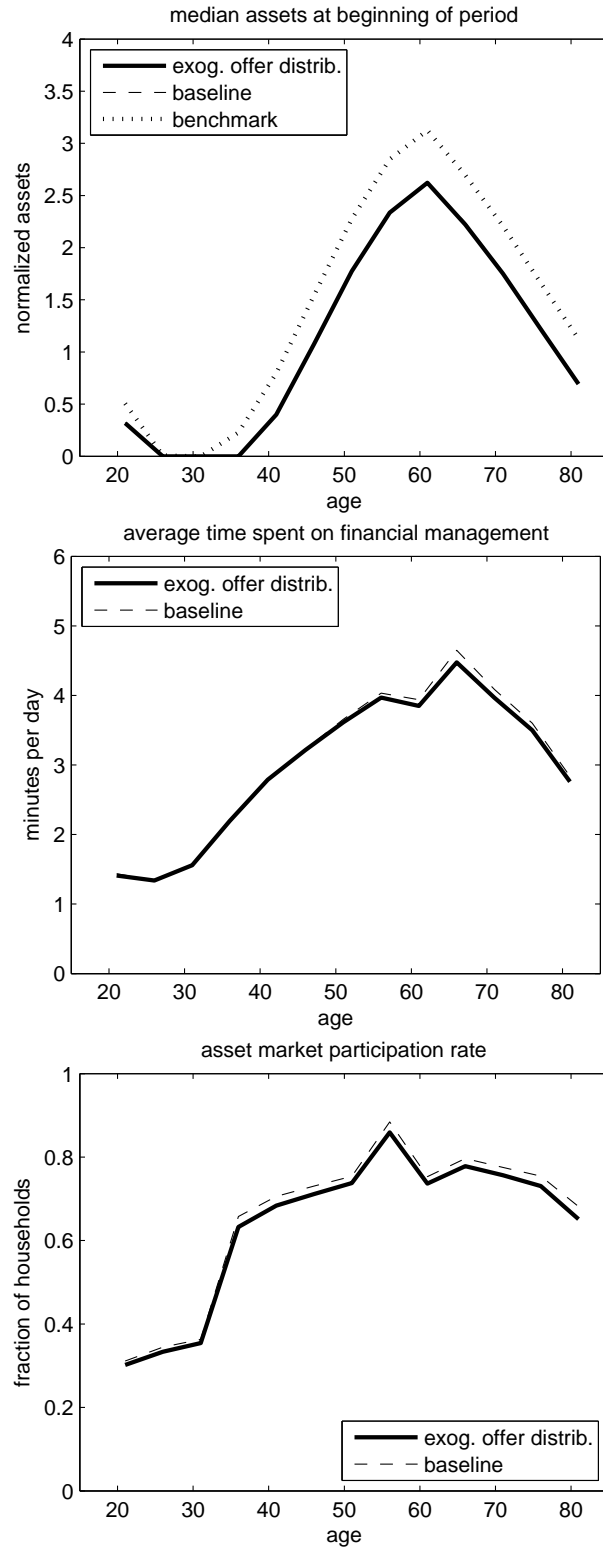


Figure 6: Life-cycle profiles of assets, search effort and asset market participation rate when offer distribution is fixed exogenously. The baseline economy has an endogenous offer distribution and the benchmark economy has a frictionless asset market.

As both the endogenous offer distribution and the fixed offer distribution provide the households with a number of attractive offers, they are close to equivalent in the eyes of households that choose moderate to high levels of search effort.

5 The privatization experiment

The results from the previous section show that the model is able to capture several features of household saving behavior that are not captured by the standard model of household saving such as limited asset market participation and dispersed returns conditional on participation. I now explore how households' financial outcomes change in response to a social security privatization reform and how asset market frictions affect the welfare consequences of social security privatization. I present results for one particular policy experiment that is similar to one conducted by Nishiyama and Smetters (2007) with a frictionless asset market.

5.1 Description of the policy experiment

Privatization is modeled as a 50% reduction in social security benefits. Social security taxes and income taxes are also reduced so that the two government budgets are balanced in the new steady state.

While modeling privatization as simply a reduction in the program is common in the literature on social security reform,²³ it has a different interpretation here. Proposals to introduce a system of private social security accounts sometimes specify that household choices be restricted to a specific set of approved assets. The experiment here implicitly assumes that the private accounts are completely unrestricted so the household is free to invest its social security savings in exactly the same way that it would invest its private savings.

In the experiment, benefits are reduced linearly over a period of 25 years and those households that are retired when the policy change is announced continue to receive the original benefit levels. The payroll tax, however, is immediately reduced to its new steady-state level, which is slightly less than 50% of its original level.

²³See Kotlikoff et al. (1998) and Nishiyama and Smetters (2007) for examples.

The pay-as-you-go system has an implicit debt to those households that have paid social security taxes but have not yet received their benefits. The gradual reduction in benefits is implemented so that those households that have already paid taxes are compensated for most of their contributions. These benefits must be funded despite the fact that payroll taxes have already been cut and there are three ways that these obligations can be dealt with: disregard them, issue new debt to pay for them or pay for them out of additional tax revenue. I choose the latter approach and impose a consumption tax on the transition generations so that the social security budget is balanced period-by-period over the transition. As Kotlikoff et al. (1998) have shown, a declining consumption tax path encourages saving and speeds up the accumulation of capital. Those authors consider consumption, income and payroll taxes and find that funding the transition with a consumption tax leads to the fastest transition.

In the initial steady state, government expenditures, Q_t , are constant in per capita terms. I assume that they remain constant in per capita terms over the course of the transition. As the economy expands after privatization, the income tax needed to finance these expenditures falls. In the experiment, the income tax rate adjusts period-by-period to maintain budget balance.

I compute a 150-year, perfect-foresight transition path from the initial steady state to the post-privatization steady state. Computing the transition requires finding a sequence of offer distributions corresponding to the firm equilibrium at each date. Appendix C describes the procedure to compute the transition. Finally, I also perform the same experiment in the benchmark model in order to understand how the search friction changes the analysis.

5.2 The long-run impact of privatization

Table 3 shows the long-run response of macroeconomic quantities to social security privatization for both the full model and the benchmark model. In both models, the capital stock expands and labor supply increases resulting in an expansion of output. The full model generates a somewhat larger response to privatization, which is consistent with the view that frictional asset markets magnify the distortions caused by social security. Nevertheless, the magnitudes of the responses are quite similar in the two models, which

	full model	benchmark model
output	4.94%	4.81%
capital	10.47%	10.26%
labor supply	3.60%	3.49%
average search effort	19.3%	–
asset market participation	11.7 %	–
mass of firms	12.0%	–

Table 3: Response to privatization experiment. The table reports the percentage difference between the initial steady-state value and the post-privatization steady-state value. Asset market participation is in percentage point difference. The benchmark model is the full model less the search friction.

may be explained by the fact that the saving behavior of wealthy households is crucial for the evolution of the aggregate capital stock and these households choose high levels of search effort both before and after privatization.

In the full model, average search effort rises by 19% and the asset market participation rate rises by 12 percentage points. This increase in participation is particularly noticeable among middle aged households. Average search effort among households nearing retirement and in early retirement rises by 30% or more. This expansion in search effort coincides with an expansion of asset market participation from roughly 80% to close to 100% for these age groups.

Privatization affects the offer distribution in two ways. The firm’s profit equation depends directly on the marginal product of capital, so firms offer lower returns as the marginal product of capital falls. The firm’s profit equation also depends on the search behavior chosen by households. A useful way of normalizing the offers to remove the direct effect of the marginal product of capital is to look at the distribution of “fees,” which I define as the difference between the offer and the marginal product of capital. Fees are generally smaller after privatization, but the difference between the two distributions is small with the median fees only differing by a few basis points. There are two reasons that this effect is so small. First, the offer distribution depends on household search behavior weighted by the household assets so the response of the wealthy households to privatization has a large role in shaping the response of the offer distribution to privatization. Wealthy households have muted responses to privatization as a large share of their income is unaffected by the change in payroll taxes and they are not particularly reliant on social security payments to

fund their consumption in retirement. Second, the fact that the average level of search effort increases does not necessarily mean that the offer distribution becomes more attractive. On the one hand, privatization causes households that already search to search harder and become more discerning, but on the other hand, privatization causes new households to enter the market that are not necessarily well informed. As these effects are partially offsetting, the offer distribution only improves slightly in response to privatization.

In the long run, privatization produces welfare gains. To compare welfare across steady states, I compute the average expected utility of a household at birth. In the full model, the long-run increase in welfare is equal to 6.1% of consumption. In the benchmark economy, the gain is larger at 6.7% of consumption.

5.3 Welfare impact including transition costs

The full benefit of privatization is not realized until many years after the reform once the transition costs have been paid. To incorporate these transition costs in the analysis, I compare the discounted expected utility of each household at the date the policy change is announced. At this date, each household's discounted expected utility changes to reflect the new policy. Table 4 presents the average change in expected utility across different dimensions of the state space. Across all agents, the policy announcement reduces welfare by 1.06% in consumption equivalents. For the benchmark economy this figure is 0.88%. That the privatization experiment produces a small welfare loss for the benchmark economy is not a surprising result. Nishiyama and Smetters (2007) conduct a similar experiment and find that privatization produces small efficiency losses when the costs of transition are taken into account.

What accounts for the difference in results across the two models? The models differ in three dimensions. First, there is the fundamental difference in the model environments, which is the search friction in the asset market. On top of this difference, the models differ in their calibrations as the benchmark model is recalibrated to fit the same set of moments. Finally, there are different general equilibrium responses to the privatization reform. To explore the roles of these three differences, I perform a partial equilibrium experiment in which I modify the welfare calculation in the benchmark economy by asking how a measure-zero group of agents with preferences calibrated as in the full model would view the reform. The answer is

		total impact				
full model		-1.06%				
benchmark model		-0.88%				
		across ages				
		21 - 25	41 - 50	56 - 65	66 - 75	81 - 95
full model		0.75%	-3.11%	-2.00%	0.46%	0.91%
benchmark model		1.03%	-2.83%	-2.11%	0.41%	1.05%
		across asset levels (percentile)				
		0 - 20	40 - 60	80 - 100		
full model		-0.90%	-0.89%	-1.41%		
benchmark model		-0.67%	-0.83%	-1.41%		
		across income levels (shock)				
		1, 2	3, 4, 5	6, 7		
full model		-2.63%	-0.98%	-0.96%		
benchmark model		-2.44%	-0.79%	-0.87%		

Table 4: Welfare impact of privatization policy announcement across different dimensions of the state space. The benchmark model is the full model less the search friction. Welfare differences are expressed in terms of consumption equivalents.

that they would realize a welfare loss of 0.84% of consumption, which is not much changed from the 0.88% for the benchmark economy. I then suppose that these agents were facing the equilibrium prices from the full model and recompute the change in welfare. Here I find a loss of 0.82% of consumption. From these experiments, I conclude that the additional welfare loss is driven by the search friction directly and not through the calibration of the model or the general equilibrium effects.

Why does the search friction lead to a larger welfare loss? Part of the benefit of the social security system is that it is a form of insurance and this insurance is reduced by the reform. The insurance value of social security is clear in the results of Nishiyama and Smetters (2007) as they find that privatization leads to a substantial efficiency gain when wage risk is insurable and a small efficiency loss when wages are not insurable. Social security also provides insurance against longevity risk in this economy in which there are no annuity markets.²⁴ The introduction of the search friction makes it more difficult for households to self insure because savings are themselves risky due to the stochastic nature of the search process. As a result, the insurance value of social security is larger with the search friction.

While the search friction does affect the details of the welfare analysis, the results are not so different

²⁴See Imrohoroglu et al. (1995) for an analysis of this role of social security.

as to lead to different conclusions about the desirability of such a reform. The take-away message is largely that the standard analysis is robust to frictions of this type.

5.4 Heterogeneous search efficiencies

To this point, I have assumed that all households are equally efficient at finding returns in the asset market. It seems relevant to ask how the analysis of social security privatization would change if some households are poorly equipped to search. Households with low search efficiencies could be made substantially worse off by privatization, which would alter the overall welfare calculations. To consider this case, I must take a stand on how search efficiency covaries with income and age. I suppose that differences in labor productivity reflect differences in human capital and human capital is also a necessary input to the search process. Therefore, I model a working household's search efficiency as proportional to their labor productivity: $\theta(s; \epsilon) = \theta_1 \times \log(1 + \theta_2 y(\epsilon)s)$, where $y(\epsilon)$ is the productivity of a household of type ϵ . When a household retires, its search efficiency remains at the final working-age level, which presumes there is no loss of ability due to cognitive decline. I recalibrate the model to match the same moments as above.

The welfare impact of privatization only changes a small amount from the case of homogeneous search efficiencies, producing a welfare loss of 1.22% rather than 1.06%. Households respond to low levels of search efficiency by increasing the amount of effort devoted to search and the cost of doing so is low as these households have low levels of labor productivity.

6 Conclusion

In this paper, I have argued that several features of the data on household savings behavior can be understood in terms of the incentives to learn about investment opportunities and I have proposed a model of savings in which households must search for returns in the asset market. The model accurately predicts that households will spend more time managing their finances and be more likely to participate in asset markets as they approach retirement. The model also predicts that wealthy and high-income households will spend

more time managing their finances, be more likely to participate in the asset market and earn higher returns conditional on participation.

The key message with regard to the social security privatization experiment is that the introduction of a friction in the asset market as modeled here makes privatization less attractive, but only slightly so. When all of the costs of the transition to the post-privatization steady state are included, the reform generates a welfare loss both with and without the frictions in the asset market. These losses, are somewhat larger with the asset market frictions, but the conclusions we draw from the standard model remain largely intact.

References

- Aiyagari, S. R., 1994. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics* 109 (3), 659–684.
- Ameriks, J., Caplin, A., Leahy, J., 2003. Wealth accumulation and the propensity to plan. *Quarterly Journal of Economics* 118, 1007–1047.
- Arrow, K. J., 1987. The demand for information and the distribution of income. *Probability in the Engineering and Informational Sciences* 1, 3–13.
- Benhabib, J., Bisin, A., Zhu, S., 01 2011. The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica* 79 (1), 123–157.
- Bewley, T., undated. Interest Bearing Money and the Equilibrium Stock of Capital. Manuscript.
- Bontemps, C., Robin, J.-M., van den Berg, G., 1997. Equilibrium search with productivity dispersion: Theory and estimation. Tinbergen Institute Discussion Paper.
- Budria Rodriguez, S., Diaz-Gimenez, J., Quadrini, V., Rios-Rull, J.-V., 2002. Updated facts on the u.s. distributions of earnings, income, and wealth. *Federal Reserve Bank of Minneapolis Quarterly Review* 26 (3), 2–35.
- Burdett, K., Judd, K. L., 07 1983. Equilibrium price dispersion. *Econometrica* 51 (4), 955–69.
- Butters, G. R., 1977. Equilibrium distributions of sales and advertising prices. *The Review of Economic Studies* 44 (3), pp. 465–491.
- Cagetti, M., De Nardi, M., 10 2006. Entrepreneurship, frictions, and wealth. *Journal of Political Economy* 114 (5), 835–870.
- Calvet, L. E., Campbell, J. Y., Sodini, P., 10 2007. Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy* 115 (5), 707–747.

- Campanale, C., 10 2007. Increasing returns to savings and wealth inequality. *Review of Economic Dynamics* 10 (4), 646–675.
- Carlin, B. I., Manso, G., 2011. Obfuscation, learning, and the evolution of investor sophistication. *Review of Financial Studies* 24 (3), 754–785.
- Castaneda, A., Diaz-Gimenez, J., Rios-Rull, J.-V., 08 2003. Accounting for the u.s. earnings and wealth inequality. *Journal of Political Economy* 111 (4), 818–57.
- Choi, J. J., Laibson, D., Madrian, B. C., 2010. Why does the law of one price fail? an experiment on index mutual funds. *Review of Financial Studies* 23 (4), 1405–1432.
- Conesa, J. C., Krueger, D., 10 1999. Social security reform with heterogeneous agents. *Review of Economic Dynamics* 2 (4), 757–795.
- De Nardi, M., 2004. Wealth inequality and intergenerational links. *Review of Economic Studies* 71 (3), 743–768.
- Diamond, P., Orszag, P., 2004. A Summary of Saving Social Security: A Balanced Approach. MIT working paper.
- Drozdz, L., Nosal, J., 2008. Competing for customers: A search model of the market for unsecured credit.
- Feldstein, M., Liebman, J. B., 2002. Social Security. *Handbook of public economics*. Volume 4. *Handbooks in Economics*, vol. 4; Amsterdam; London and New York;; Elsevier Science, North-Holland, pp. 2245–2324.
- Fuster, L., Imrohoroglu, A., Imrohoroglu, S., 2007. Elimination of social security in a dynastic framework. *Review of Economic Studies* 74 (1), 113–145.
- Gouskova, E., Chiteji, N., Stafford, F., June 2010. Estimating the intergenerational persistence of lifetime earnings with life course matching: Evidence from the psid. *Labour Economics* 17 (3), 592–597.
- Haliassos, M., Bertaut, C. C., 1995. Why do so few hold stocks? *The Economic Journal* 105 (432), pp. 1110–1129.

- Heathcote, J., Storesletten, K., Violante, G. L., 2010. The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy* 118 (4), 681 – 722.
- Hortacsu, A., Syverson, C., 05 2004. Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *Quarterly Journal of Economics* 119 (2), 403–56.
- Huang, H., Imrohoroglu, S., Sargent, T. J., January 1997. Two computations to fund social security. *Macroeconomic Dynamics* 1 (01), 7–44.
- Huggett, M., 1993. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control* 17 (5-6), 953–969.
- Huggett, M., Ventura, G., 1999. On the distributional effects of social security reform. *Review of Economic Dynamics* 2 (3), 498 – 531.
- Hurst, E., Willen, P., 08 2007. Social security and unsecured debt. *Journal of Public Economics* 91 (7-8), 1273–1297.
- Imrohoroglu, A., Imrohoroglu, S., Joines, D. H., 1995. A life cycle analysis of social security. *Economic Theory* 6, 83–114.
- Kotlikoff, L. J., Smetters, K. A., Walliser, J., 05 1998. Social security: Privatization and progressivity. *American Economic Review* 88 (2), 137–141.
- Lusardi, A., Mitchell, O. S., 01 2007. Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics* 54 (1), 205–224.
- Mankiw, N. G., Zeldes, S. P., 03 1991. The consumption of stockholders and nonstockholders. *Journal of Financial Economics* 29, 97–112.
- Mazumder, B., 07 2005. Fortunate sons: New estimates of intergenerational mobility in the united states using social security earnings data. *The Review of Economics and Statistics* 87 (2), 235–255.

- Nishiyama, S., Smetters, K., 11 2007. Does social security privatization produce efficiency gains? *Quarterly Journal of Economics* 122 (4), 1677–1719.
- Social Security Administration, 2007. Annual statistical supplement to the social security bulletin.
- Stephenson, E. F., 04 1998. Average marginal tax rates revisited. *Journal of Monetary Economics* 41 (2), 389–409.
- Storesletten, K., Telmer, C. I., Yaron, A., 1999. The risk-sharing implications of alternative social security arrangements. *Carnegie-Rochester Conference Series on Public Policy* 50 (0), 213 – 259.
- Tauchen, G., 1986. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters* 20, 177–181.
- van Rooij, M., Lusardi, A., Alessie, R., 2007. Financial Literacy and Stock Market Participation. National Bureau of Economic Research, Inc, NBER Working Papers: 13565.
- Vissing-Jorgensen, A., 2002. Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures. National Bureau of Economic Research, Inc, NBER Working Papers: 8884.
- Wade, W. R., 2000. *An Introduction to Analysis: Second Edition*. Prentice Hall, Upper Saddle River, NJ.

Appendices are not for publication

A Proofs

Proofs for section 2

Lemma 1. *Under assumption 1, if $F(r)$ is continuous in r , then*

$$p(a, \epsilon; r) \equiv \mu^{-1} \sum_{j=0}^{\infty} j q(j; a, \epsilon) [F(r)]^{j-1}$$

is continuous in r for any (a, ϵ) .

Proof. Consider two points r^1 and r^2 , and let $F(r^1) \leq F(r^2)$ without loss of generality. Let $z = F(r^2) - F(r^1)$. For any j and any z , $[F(r^2)]^j - [F(r^1)]^j$ is weakly increasing in $F(r^1)$. To see this, write the difference as

$$\begin{aligned} [F(r^2)]^j - [F(r^1)]^j &= [F(r^1) + z]^j - [F(r^1)]^j \\ &= \sum_{k=0}^j \binom{j}{k} [F(r^1)]^{(j-k)} z^k - [F(r^1)]^j \\ &= \sum_{k=1}^j \binom{j}{k} [F(r^1)]^{(j-k)} z^k, \end{aligned}$$

which is increasing in $F(r^1)$ (weakly if $z = 0$).

As $F(r)$ is continuous by assumption and since $p(a, \epsilon; r)$ only depends on r through $F(r)$ it is sufficient to prove that p is continuous in F for all $F \in [0, 1]$. Fix an $\epsilon > 0$. Assumption 1 implies that there is a J such that $\sum_{j=J}^{\infty} j q(j; a, \epsilon) < \epsilon / (2\mu)$. Therefore, for any $F(r^1) \leq F(r^2) \in [0, 1]$,

$$p(a, \epsilon; r^2) - p(a, \epsilon; r^1) < \mu \sum_{j=0}^{J-1} j q(j; a, \epsilon) \left\{ [F(r^2)]^{j-1} - [F(r^1)]^{j-1} \right\} + \epsilon/2.$$

Following the discussion above, for a given $z = F(r^2) - F(r^1)$, the difference $[F(r^2)]^{j-1} - [F(r^1)]^{j-1}$ is

largest for $F(r^2) = 1$. Moreover, $1 - [F(r^1)]^j$ is weakly increasing in j . Therefore,

$$p(a, \epsilon; r^2) - p(a, \epsilon; r^1) < \mu \left[1 - (1 - z)^{J-2} \right] \sum_{j=0}^{J-1} jq(j; a, \epsilon) + \epsilon/2.$$

As all q 's are positive, assumption 1 implies $\sum_{j=0}^{J-1} jq(j; a, \epsilon)$ is finite. By choosing $z = F(r^2) - F(r^1)$ to be sufficiently small, the first term in the expression above can be made less than $\epsilon/2$. Hence, p is continuous in F . ■

The following proof is similar to that of lemma 1 of Burdett and Judd (1983).

Proof of proposition 1. $F(\cdot)$ is continuous. Suppose to the contrary that there is a mass of firms offering the same return r . Then any one of those firms could profitably deviate to $r + \epsilon$ for some sufficiently small $\epsilon > 0$. To see this, write the profits as

$$\begin{aligned} \pi(r) &= \left(1 - \frac{1+r+\epsilon}{A} \right) \sum_{\epsilon} \int h(a, \epsilon) p(a, \epsilon; r) \Phi(da, \epsilon) \\ p(a, \epsilon; r + \epsilon) &= \mu \sum_{j=0}^{\infty} jq(j; a, \epsilon) [F(r + \epsilon)]^{j-1}. \end{aligned}$$

Note that as F is discontinuous at r , $p(a, \epsilon; r + \epsilon)$ is discretely larger than $p(a, \epsilon; r)$ at all (a, ϵ) , but $[1 - (1 + r + \epsilon)(A)]$ is only lower than $[1 - (1 + r)(A)]$ by the arbitrarily small amount ϵ/A so there is a sufficiently small ϵ for which $r + \epsilon$ represents a profitable deviation from r .

The support of $F(\cdot)$ starts at zero. An offer below the reservation return would not attract any investments so the expected profit (before the fixed cost) is zero while offering a return of zero (the reservation return) would yield a positive expected profit as, by assumption 2(i), some households only encounter a single firm. Thus, negative returns are dominated so the support of the offer distribution does not begin below zero. Now consider a firm making an offer r for which $F(r) = 0$. Such a firm only receives investments from households who have encountered no other firms so the firm can offer the reservation return and receive no fewer investments.

The support of $F(\cdot)$ ends at some $\bar{r} < A - 1$. Suppose to the contrary that $\bar{r} \geq A - 1$. Firms offering

$A - 1$ or more would make zero profits at best and a deviation to offering the reservation return would be profitable.

$F(\cdot)$ is strictly increasing on $[0, \bar{r}]$. Consider two points $r^1 < r^2 \in [0, \bar{r}]$. Suppose to the contrary that $F(r^1) = F(r^2)$. If $F(r^1) = 1$ then $r^1 \geq \bar{r}$ and $r^2 > \bar{r}$. Thus, $F(r^1) < 1$. As $F(\cdot)$ is continuous, for any $\varepsilon > 0$, there must be some $\tilde{r}^2 \geq r^2$ such that $F(\tilde{r}^2) < F(r^1) + \varepsilon$ and \tilde{r}^2 is offered in equilibrium. As $p(a, \varepsilon; r)$ is continuous in r , for any $\varepsilon' > 0$ there is an ε such that $p(a, \varepsilon; \tilde{r}^2) - p(a, \varepsilon; r^1) < \varepsilon'$. Therefore, using the notation $S(r) = \pi(r) \times A / (A - 1 - r)$,

$$\begin{aligned} S(\tilde{r}^2) - S(r^1) &= \sum_{\varepsilon} \int h(a, \varepsilon) [p(a, \varepsilon; \tilde{r}^2) - p(a, \varepsilon; r^1)] \Phi(da, \varepsilon) \\ &< \varepsilon' \sum_{\varepsilon} \int h(a, \varepsilon) \Phi(da, \varepsilon). \end{aligned}$$

Assumption 2(iii) implies the integral is a finite constant so this difference can be made arbitrarily small through an appropriate choice of ε' and ε . Now consider the profit from offering \tilde{r}^2 rather than r^1 .

$$\begin{aligned} \pi(\tilde{r}^2) - \pi(r^1) &= \left(1 - \frac{1 + \tilde{r}^2}{A}\right) S(\tilde{r}^2) - \left(1 - \frac{1 + r^1}{A}\right) S(r^1) \\ &= \left(1 - \frac{1 + \tilde{r}^2}{A}\right) [S(\tilde{r}^2) - S(r^1)] + \frac{r^1 - \tilde{r}^2}{A} S(r^1). \end{aligned}$$

For sufficiently small ε , the first term in the expression above can be made arbitrarily small while the second is fixed and negative as $S(r^1)$ is positive by assumption 2(i). This implies for small ε , $\pi(r^1) > \pi(\tilde{r}^2)$ so \tilde{r}^2 is not offered in equilibrium, but \tilde{r}^2 is offered by construction, hence a contradiction. ■

B Heterogeneous firms

It is assumed above that all firms have marginal product of capital A . The goal of this appendix is to show that the results are robust to firm heterogeneity. To accomplish this, this appendix considers a sequence of firm equilibria in which the firms' marginal products are distributed on the sequence of intervals $[\underline{A}_n, \bar{A}_n]$.

The main result shows that if $\underline{A}_n \rightarrow A$ and $\bar{A}_n \rightarrow A$ for some A then the associated sequence of offer distributions converges to the one that arises in the case in which firms are homogeneous with marginal product of capital A . In this appendix the behavior of households is taken as given, which implies that there is a unique firm equilibrium. In addition, the behavior of households is assumed to satisfy assumption 2.

Let $\{D_n(A)\}_{n=1}^\infty$ be a sequence of distributions of marginal products with associated supports $\{[\underline{A}_n, \bar{A}_n]\}_{n=1}^\infty$ that satisfy $\underline{A}_n \rightarrow A$ and $\bar{A}_n \rightarrow A$ as $n \rightarrow \infty$. Let $\{F_n(r)\}_{n=1}^\infty$ be the sequence of offer distributions associated with the sequence of marginal product distributions $\{D_n(A)\}_{n=1}^\infty$. The goal is to show that $\{F_n(r)\}_{n=1}^\infty \rightarrow F(r)$ point-wise as $n \rightarrow \infty$, where $F(r)$ is an offer distribution that arises in the equilibrium with homogeneous firms.

As a matter of notation, note that equation 2 can be rewritten as

$$\pi_n(r, A) = \left(1 - \frac{1+r}{A}\right) S_n(r),$$

where $S_n(r) = \int \int h(a, \epsilon) p_n(a, \epsilon; r) \Phi(da, d\epsilon)$ represents the expected investments when a return r is offered, which is independent of the marginal product of capital. $p_n(a, \epsilon; r)$ is given by

$$p_n(a, \epsilon; r) = \left\{ \mu^{-1} \sum_{j=0}^{\infty} j q(j; a, \epsilon) [F_n(r)]^{j-1} \right\}.$$

Notice that $S_n(r)$ depends on the offer distribution $F_n(r)$ through $p_n(a, \epsilon; r)$. $S(r)$ will refer to the same object in the homogeneous case. In addition, let $\bar{r}_n = \sup \{r : F_n(r) < 1\}$ and for the homogeneous case let $\bar{r} = \sup \{r : F(r) < 1\}$.

Before stating and proving this proposition, a few lemmas are necessary.

Lemma 2. *For any n , $F_n(r)$ is continuous and strictly increasing with $\inf \{r : F(r) > 0\} = 0$.*

Proof. The proof is the same as for the homogeneous case. ■

Lemma 3. *$S_n(r)$ is continuous for any n .*

Proof. This follows from lemmas 1 and 2. ■

The fact that $F_n(r)$ is strictly increasing does not necessarily imply that all points in the interval $[0, \bar{r}_n]$ are offered in equilibrium as a single point of measure zero could be omitted from the support of $F_n(r)$ without difficulty. This possibility motivates the next lemma, which states that the support of $F_n(r)$ is dense in $[0, \bar{r}_n]$.

Lemma 4. *For any n and any $\varepsilon > 0$, if there are three points x , \underline{x} and \bar{x} such that \underline{x} and \bar{x} are in the support of $F_n(r)$ and $\underline{x} \leq x \leq \bar{x}$, then there is a point y in the support of $F_n(r)$ such that $|x - y| < \varepsilon$.*

Proof. If $|x - \bar{x}| < \varepsilon$ or $|x - \underline{x}| < \varepsilon$ the result is immediate. Otherwise, suppose to the contrary that there were points x , \underline{x} and \bar{x} such that \underline{x} and \bar{x} are in the support of $F_n(r)$ and $\underline{x} \leq x \leq \bar{x}$, but there is no such point y in the interval $(x - \varepsilon, x + \varepsilon)$. Let $\bar{x}' = \sup \{r : F_n(r) = F_n(x)\}$ and $\underline{x}' = \inf \{r : F_n(r) = F_n(x)\}$. By assumption $\bar{x}' - \underline{x}' > 2\varepsilon$. By construction there are points in the support of F_n that are arbitrarily close to \bar{x}' and \underline{x}' . As $F_n(\bar{x}') = F_n(\underline{x}')$ it follows that profits from offering \underline{x}' are strictly greater than those from offering \bar{x}' for any marginal product A , which, by the continuity of $F_n(r)$ and $\pi(r, A)$, is inconsistent with points in a neighborhood around \bar{x}' being in the support of $F_n(r)$. ■

Lemma 5. *For any n and x , if there are points $\bar{x} \geq x$ and $\underline{x} \leq x$ such that \bar{x} and \underline{x} are in the support of F_n then $S_n(x)$ must satisfy the following inequalities*

$$\frac{\bar{A}_n - 1}{\bar{A}_n - 1 - x} \leq \frac{S_n(x)}{S_n(0)} \tag{6}$$

$$\frac{\underline{A}_n - 1}{\underline{A}_n - 1 - x} \geq \frac{S_n(x)}{S_n(0)}. \tag{7}$$

Proof. There are two cases to consider: either x is in the support of F_n or $\underline{x} < x < \bar{x}$ and x is not in the support of F_n . Suppose x is in the support of F_n . Let A^1 denote the marginal product of the firm that offers x . This firm must not wish to deviate to offer a return of 0. That is

$$(A^1 - 1) S_n(0) \leq (A^1 - 1 - x) S_n(x),$$

which can be rearranged as

$$\frac{A^1 - 1}{A^1 - 1 - x} \leq \frac{S_n(x)}{S_n(0)}.$$

And $A^1 \leq \bar{A}_n$ implies inequality (6). Analogous steps lead to inequality (7).

Now consider the case in which x is not in the support of F_n and $\underline{x} < x < \bar{x}$. By using lemma 4 repeatedly, it is possible to construct a sequence $\{x_i\}_{i=1}^{\infty}$ that converges to x with the property that every element of the sequence is in the support of F_n . The argument above applies to all such points. By the Comparison Theorem for Functions,²⁵ it follows that $\lim_{x_i \rightarrow x} S_n(x_i)$ satisfies the inequalities. As $S_n(\cdot)$ is continuous, $S_n(x) = \lim_{x_i \rightarrow x} S_n(x_i)$. ■

Lemma 6. $\bar{r}_n \rightarrow \bar{r}$ as $n \rightarrow \infty$.

Proof. In the homogeneous case, firms are indifferent between offering all returns and in particular the highest and lowest returns so \bar{r} must satisfy the condition

$$(A - 1) S(0) = (A - 1 - \bar{r}) S(\bar{r}). \quad (8)$$

As $F(\bar{r}) = 1$, $S(\bar{r})$ is given by

$$S(\bar{r}) = \mu \int \int h(a, \epsilon) \sum_{j=0}^{\infty} j q(j; a, \epsilon) \Phi(da, d\epsilon), \quad (9)$$

which is independent of the offer distribution. Similarly, as $F(0) = 0$, $S(0)$ is also independent of the offer distribution and is given by

$$S(0) = \mu \int \int h(a, \epsilon) q(1; a, \epsilon) \Phi(da, d\epsilon). \quad (10)$$

²⁵See Wade (2000), p. 62.

Finally, equation (8) can be rearranged as

$$\bar{r} = A - 1 - (A - 1) \frac{S(0)}{S(\bar{r})}, \quad (11)$$

where $S(0)$ and $S(\bar{r})$ are given by the expressions above.

The proof now proceeds by showing that for any $\varepsilon > 0$ there is an N such that $n > N$ implies the support of F_n lies below $\bar{r} + \varepsilon$ and there is a point x in the support of F_n that satisfies $|x - \bar{r}| < \varepsilon$. Consider a firm that offers $\bar{r} + \varepsilon$ or more. Such a firm must not have an incentive to deviate to offering a return of zero. That is

$$(A^1 - 1 - \bar{r} - \varepsilon) S_n(\bar{r} + \varepsilon) \geq (A^1 - 1) S_n(0)$$

for some $A^1 \in [\underline{A}_n, \bar{A}_n]$. As $F_n(\bar{r} + \varepsilon) \leq 1$, it must be the case that $S_n(\bar{r} + \varepsilon) \leq S(\bar{r})$ as given in equation (9). Moreover, it is still the case that $F_n(0) = 0$ so $S_n(0) = S(0)$ as given in equation (10). Making these substitutions, substituting equation (11) for \bar{r} and rearranging yields

$$(A^1 - A) \left(1 - \frac{S(0)}{S(\bar{r})} \right) \geq \varepsilon.$$

The first term on the left-hand side goes to zero as $n \rightarrow \infty$ and the second term is a positive constant. Therefore, this condition cannot hold for any $\varepsilon > 0$ as $n \rightarrow \infty$. Thus, $\lim_{n \rightarrow \infty} \bar{r}_n \leq \bar{r}$.

Suppose that the support of F_n lies entirely below $\bar{r} - \varepsilon$ for some $\varepsilon > 0$. Then $S_n(\bar{r} - \varepsilon) = S(\bar{r})$. Moreover, it is still the case that $S_n(0) = S(0)$. Now consider a firm that offers a return of arbitrarily close to zero (by lemma 2 there must be such a firm). This firm must not have an incentive to deviate to offering \bar{r} , which any firm is free to do. That is

$$(A^1 - 1) S(0) \geq (A^1 - 1 - \bar{r} + \varepsilon) S(\bar{r})$$

inserting equation (11) for \bar{r} and rearranging yields

$$(A^1 - A) \left(\frac{S(0)}{S(\bar{r})} - 1 \right) \geq \varepsilon.$$

The left-hand side of this expression goes to zero as $n \rightarrow \infty$ so the condition does not hold in the limit.

Thus, $\lim_{n \rightarrow \infty} \bar{r}_n \geq \bar{r}$ and so $\lim_{n \rightarrow \infty} \bar{r}_n = \bar{r}$. ■

Proposition 2. *Let $\{D_n(A)\}_{n=1}^{\infty}$ be a sequence of distributions of marginal products with associated supports in $\{[\underline{A}_n, \bar{A}_n]\}_{n=1}^{\infty}$. For each $D_n(A)$, there is an offer distribution $F_n(r)$. As $\underline{A}_n \rightarrow A$ and $\bar{A}_n \rightarrow A$, $F_n(r) \rightarrow F(r)$, where $F(r)$ is the offer distribution that arises when all firms are homogenous with marginal product A .*

Proof. Fix a point $x \in (0, \bar{r})$. By lemma 6 there is an N such that for $n > N$ there are points \underline{x}_n and \bar{x}_n in the support of F_n such that $\underline{x}_n \leq x \leq \bar{x}_n$. Therefore lemma 5 applies and for $n > N$ it is the case that

$$\frac{\bar{A}_n - 1}{\bar{A}_n - 1 - x} \leq \frac{S_n(x)}{S_n(0)} \leq \frac{\underline{A}_n - 1}{\underline{A}_n - 1 - x}$$

Recall that $S_n(0) = S(0)$ for all n , so by the Squeeze Theorem it follows that

$$\lim_{n \rightarrow \infty} S_n(x) = \frac{A - 1}{A - 1 - x} S(0). \quad (12)$$

From the equal-profit condition in the homogeneous case we have

$$(A - 1 - x) S(x) = (A - 1) S(0),$$

so equation (12) implies $S_n(x) \rightarrow S(x)$.

Notice that $S_n(\cdot)$ depends on n only through F_n so it may be written as $S(F_n(\cdot))$. Given assumption 2, $S(\cdot)$ is continuous and strictly increasing in its argument $F \in [0, 1]$ and is therefore invertible. As a result, convergence of $S_n(x)$ to $S(x)$ implies convergence of $F_n(x)$ to $F(x)$. ■

C Computational details

C.1 Computing the steady state

For a given set of parameter values, the algorithm begins with a guess of the capital-labor ratio and the associated wage and marginal product of capital. Next I guess a distribution of offered returns. The consumer's problem can then be solved and simulated. After solving the household's problem, the next step is to simulate a sample of households from which one can compute the corresponding firm equilibrium offer distribution from the firm's profit equation. I then solve the household problem again until the offer distribution converges (it usually does so after a small number of iterations). Finally, I check the simulated capital-labor ratio and updates the guess.

To solve the household's problem I use a value function iteration algorithm in which the value function for each discrete type is interpolated over 35 unequally spaced nodes using cubic splines. Solving the household's problem requires integrating over returns with respect to $G(r, s)$. To compute these integrals I use Chebyshev quadrature with 21 nodes over $(0, \bar{r})$ and an additional node at $r = 0$. To compute the density of $G(r, s)$ at each node, I differentiate equation (5), which gives the density as a function of s , $F(r)$ and the density of $F(r)$. The latter two objects are stored when I solve the firm's problem.

Simulating the household's problem requires sampling from the distribution $G(r, s)$. I do this by approximating $F(r)$ with a discrete grid of points from which equation (5) gives a discrete approximation to $G(r, s)$ for a given level of search effort. I then draw from a uniform distribution and numerically invert $G(r, s)$ using this discrete approximation.

Computing the firm equilibrium simply requires solving a differential equation. As firms are indifferent between offered returns, the derivative of the firm's profit equation with respect to r must be zero on the support of $F(r)$. Rearranging this derivative produces a differential equation that $F(r)$ must satisfy. Proposition 1 provides an initial condition of $F(0) = 0$ that can be used to solve for $F(r)$.

age group	21 – 25	26 – 30	31 – 35	36 – 40	41 – 45	46 – 50	51 – 55	56 – 60	61 – 65
\bar{y}_i	1.00	1.28	1.45	1.56	1.62	1.63	1.59	1.53	1.38

Table 5: Life-cycle profile in labor productivity.

C.2 Computing the transition

The computational procedure used to find the transition path again resembles the one would use to compute the transition of a standard Bewley-Huggett-Aiyagari economy. First, one chooses a number of periods after which the economy is assumed to be in the new steady state. Next one guesses on a path for the capital-labor ratio. Given the capital-labor ratio, one guesses a sequence of offer distributions. One then solves the consumer problem backwards from the new steady state to find the households' decision rules and then uses those decision rules to simulate forward from the initial distribution of households over the state space. Using the simulated sample, one can compute solve the firm's problem for a new sequence of offer distributions. The algorithm updates the offer distributions and iterates until they converge while holding the capital-labor ratio fixed. Once this has been accomplished one compares the simulated capital-labor ratio to the guess and updates the guess. This procedure is repeated until the simulated capital labor ratio matches the guess. An outer loop checks that the sequence of consumption taxes clears the social security budget and updates the tax sequence if it does not. Finally, one must check that the economy does in fact reach the new steady state in the number of periods assumed at the outset.

D Estimating the life-cycle profile of labor productivity

To estimate the deterministic life-cycle profile in labor productivity I use data from the 1968 - 2005 waves of the PSID. I create a measure of the household wage by summing the total labor income of the head and wife and dividing by the total hours of the head and wife. I require that households work at least 208 hours to be included in the sample and that the calculated wage be at least one-half of the minimum wage in the year in question. Households are divided into nine age groups 21-25, ..., 61-66 and I regress the log wage on year and age group dummies. The resulting life-cycle profile for wages appears below.

