What Drives Investors' Portfolio Choices? Separating Risk Preferences from Frictions

Taha Choukhmane* Tim de Silva*

February 10, 2022

Preliminary and incomplete. Please do not cite or circulate.

Abstract

We separately identify the role of risk preferences and frictions in portfolio choice. Standard cross-sectional moments, such as the life cycle of participation in the stock market, cannot distinguish between the two: agents may choose not to participate because of non-standard preferences (e.g. loss aversion) or frictions impacting their choices (e.g. inertia or participation costs). We overcome this identification problem by using variation in the default asset allocation of 401(k) plans and estimate that, absent frictions, 95% of investors would prefer stock market participation with a stock share that declines over the life cycle, which differs markedly from their observed choices. We use this variation to structurally estimate a life cycle portfolio choice model and find the evidence consistent with a relative risk aversion of 3.6. This estimate is significantly lower than most estimates in the life cycle portfolio choice literature and highlights how choice frictions can hamper the identification of risk preferences.

*MIT Sloan. First draft: November 2021. We thank audience members at MIT and at Duke Fuqua for their insightful comments. We thank the data provider, a large U.S. financial institution, for making available the data used in this paper, helpful discussions, and technical support.

Many households do not participate in the stock market, including households with significant financial wealth (Mankiw and Zeldes 1991; Guiso, Haliassos, and Jappelli 2002; Campbell 2006). This limited participation in the stock market is difficult to reconcile with standard economic theory, which predicts that all investors should hold at least a small amount of stocks in the presence of a positive equity premium (e.g. Merton 1969; Campbell and Viceira 2001).[1] In principle, an investor may choose to not allocate their financial wealth to the stock market for two reasons. First, this investor may prefer holding safe assets over stocks (e.g. due to loss aversion, ambiguity aversion or pessimistic beliefs about returns). Alternatively, this investor may prefer holding stocks over safe assets, but not participate due to frictions. These frictions could include the real costs of setting up and maintaining a brokerage account or the cognitive cost of making a financial plan and paying attention. Although these two explanations have similar predictions for this investor's participation in the stock market, distinguishing between them has important normative implications. For example, interventions designed to encourage more stock market participation may be more desirable if non-participation is due to high participation costs rather than a preference for safe assets.

In this paper, we propose and implement a new empirical approach to recover investor preferences in the presence of frictions and, by doing so, distinguish between the two leading classes of explanations for limited stock market participation. We begin by showing that these two classes of explanations are difficult to separate in commonly-used cross-sectional data. Using a standard life cycle portfolio choice model, we find the life cycle profile of participation and equity shares in the Survey of Consumer Finances is consistent with two very different calibrations of the model: (i) a risk aversion below 2 and an extremely high adjustment or participation cost and (ii) a risk-aversion above 30 and a lower adjustment or participation cost. This finding illustrates the challenge in separately identifying investors' risk preferences and choice frictions, highlighting the potential limitations of previous research that has largely relied on such cross-sectional moments to estimate life cycle portfolio choice models.

Next, we overcome this identification problem using quasi-experimental variation in the default asset allocation of 401(k) plans. An ideal experiment for separating between preference- and friction-based explanations for non-participation would be to randomly give investors who are not participating in the stock market an investment account with stocks, thereby removing any one-time costs associated with participation. If these investors dislike holding risky assets (for instance, due to loss aversion) they should sell the stocks and move their holdings toward safer assets. Alternatively, if frictions were responsible for these investors not participating beforehand,

[1] The intuition is that with strictly increasing and differentiable utility, agents should be risk-neutral over small risk (Rabin 2000).

we would expect them to keep the stocks, durably switching from stock-market non-participation to participation as a consequence of the treatment.

Using account-level data from a large 401(k) plan provider, we exploit changes in the default asset allocation of retirement plans to approximate this ideal experiment. Our treatment group consists of employees hired right after the default asset allocation was changed to a Target Date Fund (TDF), which has significant equity exposure. By default, these employees are participating in the stock market, but can chose to opt-out and move their savings toward safer assets. We consider two control groups of employees hired right before the policy change: (i) employees automatically enrolled into a money market fund (with no equity exposure); (ii) employees hired under an opt-in regime. Employees in both control groups are, by default, not holding any stocks in their retirement account and must make an active decision in order to participate.[2] We find that more than 90% of employees defaulted into stocks maintain a positive equity exposure throughout their tenure, whereas employees defaulted into a money market fund (or hired under an opt-in regime) progressively move away from their zero-stock default option and toward higher equity exposure. The fact that most investors move away from the default when it is a safe asset but stay invested in the default when it's equity suggests, absent frictions, these investors prefer holding risky assets.

To translate this variation into estimates of preferences, we apply a framework developed by Goldin and Reck (2020). Under a set of assumptions, most importantly that treatment is randomly assigned and investors deviating from their default asset allocation reveal their preferences, we can non-parametrically bound the fraction of investors who would prefer stock market participation absent frictions. Empirically, 42% of investors who were defaulted into a money market opted out of the default within one year and moved toward stock market participation. Intuitively, these investors have revealed their preference for stocks by actively moving away of the money market default and, therefore, there must be at least 42% of investors who prefer stock market participation. Conversely, 5% of investors who were defaulted into holding stocks made an active decision to opt-out of stock market participation, thus revealing their preference for non-participation. Thus, the upper bound for stock market participation absent any friction should be 95%. We can also bound the average preferred stock share between 39% and 80%.

Under stronger assumptions, we can non-parametrically obtain point estimates for preferences absent frictions: we find that 94% of investors in our sample prefer stock market participation with an average preferred stock share of 75%. Moreover, we estimate preferences for participation are

[2] A natural concern is that investors may hold stocks outside of their retirement account. However, according to the Survey of Consumer Finances in 2007-2016 (our sample period), less than 5% of households participate in the stock market only outside their retirement accounts.

flat over the life cycle, while the average preferred stock share increases slightly between ages 20 and 30, but decreases after 30 at a level and rate roughly consistent with textbook theory (e.g. Merton 1969). Crucially, our estimates of *preferences* differ substantially from observed *choices* – average participation and stock shares are substantially lower and increase over the life cycle.

We view these new quasi-experimental moments as better calibration or estimation targets for both standard and behavioral models of portfolio choice. On the one hand, models abstracting from frictions could directly target our friction-free estimates of participation and equity share over the life cycle to obtain better identified preference estimates. On the other hand, richer models featuring choice frictions could exploit the difference in behavior across the treatment and control group to separately identify the size of the friction from preference estimates. To illustrate these points, we build a rich life cycle portfolio choice model and estimate it using the variation from our quasi-experiments. This model has three key features that are required to match our quasi-experimental evidence. First, investors can choose different asset allocations for existing wealth and new contributions in their retirement account. Second, investors are subject to default effects when making portfolio choice and savings decisions, where the defaults are employer-specified. Finally, investors face uncertainty about future earnings and employment status at their current employer that creates value for investors to delaying adjustments in portfolio and savings decisions away from their defaults.

Within this model, we conduct an experiment analogous to the first quasi-experiment in our data by randomizing the default asset allocation of investors when they change jobs. We then structurally estimate investors' preferences and the magnitude of the frictions they face by matching the results of the experiment in the model to our empirical results. In our estimation, we find the model closely matches our quasi-experimental evidence with a coefficient of relative risk aversion of 3.6 and a portfolio adjustment cost of $450. This estimate of risk aversion is lower than most estimates in the life cycle portfolio choice literature. A notable exception is Briggs, Cesarini, Lindqvist, and Östling (2021), who use quasi-experimental variation from lottery-winnings in Sweden and estimate a relative risk aversion of around 2.7. Collectively, these results highlight how frictions distort the mapping between observed portfolio choices and risk preferences, and how quasi-experimental variation can help overcome this identification problem.

In ongoing work, we use this model to quantify the welfare effects of target date funds. An important feature of TDFs is their automatic re-balancing, which is valuable to investors who face frictions associated with portfolio adjustment. Using our structural estimates of risk preferences and frictions, we can decompose the welfare effects of TDFs into the gains from automatic re-balancing and any losses from a possibly sub-optimal glide path. Additionally, we use the model

to study the optimal design of a default asset allocation, given the risk preferences and frictions we estimate.

Related literature. This paper makes a number of contributions relative to previous literature. First, this paper adds to the extensive literature on limited household stock market participation by offering a way to distinguish between the two leading classes of explanations for this fact. The first category, which we call "preference-based" explanations, argues investors prefer holding safe over risky assets for reasons not captured by standard models (e.g. due to non-standard preferences or beliefs). For example, investors' preferences might exhibit first-order risk aversion that makes risky assets unattractive, despite their positive expected return. This occurs in theories of loss-aversion with respect to wealth or news (Gomes 2005; Pagel 2018), narrow-framing (Barberis, Huang, and Thaler 2006), rank-dependence (Chapman and Polkovnichenko 2009), disappointment-aversion (Chapman and Polkovnichenko 2009), or ambiguity-aversion (Epstein and Wang 1994). In addition, households may perceive risky assets to have a less attractive return due to background risk (Benzoni, Collin-Dufresne, and Goldstein 2007; Catherine 2020), disaster risk (Fagereng, Gottlieb, and Guiso 2017), overly pessimistic beliefs (Briggs et al. 2021), or lack of trust in the financial sector (Guiso, Sapienza, and Zingales 2008).

The second category of explanations for limited stock market participation, which we call "friction-based", argues households prefer risky assets, but do not invest in them because of frictions associated with doing so. Such frictions could be one-time participation costs (Haliassos and Michaelides 2003; Gomes and Michaelides 2005; Abel, Eberly, and Panageas 2013) or per-period participation costs (Vissing-Jørgensen 2002; Fagereng et al. 2017; Briggs et al. 2021). These costs could be real, such as the cost of opening a brokerage account or paying a financial advisor, or psychological, such as the hassle cost of deviating from a default asset allocation or engaging in financial planning. The results in this paper support this second category of explanations for limited stock market participation based on choice frictions over explanations based on preferences and beliefs. Within the class of models with frictions, we find support for models with one-time participation or fixed adjustment costs as opposed to models with per-period participation costs. Indeed, in the presence of large enough per-period participation costs workers who are automatically enrolled into a stock fund should opt-out and move their savings away from equity, which is not what we observe in the data.

This paper also contributes to the literature on life cycle portfolio choice, initiated by Merton (1969) and surveyed by Campbell and Viceira (2001), Gomes (2020), and Gomes, Haliassos, and Ramadorai (2020). With standard preferences and processes for income and asset returns (e.g. Cocco, Gomes, and Maenhout 2005), this literature finds it is difficult to match the cross-sectional

averages of households' portfolio choices, such as participation or stock shares over the life cycle. Our results suggest this discrepancy is likely driven by frictions associated with stock market participation rather than different forms of preferences, beliefs, or risks omitted from standard models.

Additionally, this paper contributes to a growing literature on target date funds (e.g. Parker, Schoar, and Sun 2020; Duarte, Fonseca, Parker, and Goodman 2021; Gomes, Michaelides, and Zhang 2021; Massa, Moussawi, and Simonov 2021). By combining quasi-experimental variation with a life cycle model, we directly estimate the size of the frictions associated with portfolio adjustments that household face. Identifying the size of these frictions allows us to quantify the total welfare gains associated with TDFs, including the benefits associated with automatic rebalancing. Our estimate of non-trivial frictions is consistent with existing evidence of inertia in households' portfolio rebalancing (Brunnermeier and Nagel 2008; Calvet, Campbell, and Sodini 2009) and mortgage refinancing (Andersen, Campbell, Nielsen, and Ramadorai 2020; Laibson, Maxted, and Moll 2020).

Finally, our paper contributes more broadly to the literature on behavioral welfare economics (e.g. Bernheim and Rangel 2009; Allcott and Taubinsky 2015). Most related is the framework we apply for inferring preferences in the presence of framing effects developed by Goldin and Reck (2020). Our results caution against structurally estimating behavioral models using cross-sectional moments, as choice frictions can break the mapping from observed choices to underlying preferences. More constructively, this paper illustrates how quasi-experimental variation can provide useful variation for distinguishing between different behavioral theories (see also Briggs et al. 2021; Choukhmane 2021; Reher and Sokolinski 2021).

Outline. Section 1 illustrates the difficulty in separating preferences and frictions using standard cross-sectional moments. Section 2 describes our data and quasi-experimental variation, while Section 3 uses this variation to non-parametrically estimate investors' risk preferences. Section 4 describes our life cycle model and the results from a preliminary calibration. Section 5 concludes. Figure and Tables are presented at the end of the text, followed by an Appendix containing additional results, model details, and derivations.

1 Preferences versus Frictions

In this section, we discuss how cross-sectional moments cannot separately identify risk preferences and frictions. We first discuss estimates of preferences and frictions in the existing literature

on life cycle portfolio choice, highlighting how estimates of preferences vary based on the assumptions about frictions. Next we show the difficulty in separating the two in a standard life cycle portfolio choice model.

1.1 Risk Preferences in Life Cycle Portfolio Choice

A large literature on life-cycle portfolio choice numerically solves dynamic life-cycle portfolio choice models with risky labor income. Cocco et al. (2005) present a canonical model with a calibrated income process and preference parameters and show the model predicts a high equity share that is declining over the life cycle. These predictions are at odds with most data on households portfolio choices, which find a relatively low equity share that is flat or decreasing over the life cycle (e.g. Ameriks and Zeldes 2004).

Because of a discrepancy between the theoretical predictions and the data, a large literature (summarized by Gomes 2020) has proposed extensions of the model in Cocco et al. (2005), in which households either (i) have different risk preferences or beliefs or (ii) face frictions associated with participating in the stock market. Broadly speaking, the typical approach in this literature is to propose a theoretically plausible mechanism for this discrepancy and then estimate the strength of this mechanism by matching cross-sectional moments from the model to the data, such as the life cycles of participation or equity shares.

In Figure 1, we show the values of relative risk aversion calibrated or estimated from papers that solve a dynamic life-cycle portfolio choice model with standard preferences.[3] The first takeaway from Figure 1 is the substantial variation in estimates of relative risk aversion in life cycle portfolio choice models, from around 3 in Briggs et al. (2021) to around 14 in Fagereng et al. (2017) and Dahlquist, Setty, and Vestman (2018). Given these papers generally target the same cross-sectional moments, this variation in estimates of risk preferences depending on the assumptions about frictions suggests these moments are having trouble distinguishing between the two (Andrews and Shapiro 2020).

The second takeaway from Figure 1 is the average value of relative risk aversion in Figure 1 is around 8, which is much higher than estimates from other literatures. For example, the literature on life cycle consumption-savings decisions finds estimates around 1-2 (Gourinchas and Parker 2002) and the evidence from labor supply elasticities suggests an estimate around 1 (Chetty 2006).

[3] We only include papers that use time-separable CRRA or Epstein-Zin preferences. Since these papers are generally focused on matching portfolio choice rather than wealth accumulation decisions, we just focus on the coefficient of relative risk aversion, even though it also controls the EIS in the CRRA case.

The fact that relatively high risk aversion is needed to match portfolio choices is the flip side of the equity premium puzzle (Hansen and Singleton 1982; Mehra and Prescott 1985).

In sum, evidence from existing literature suggests separately identifying preferences and frictions is challenging using cross-sectional moments. In the next section, we formalize this point using a simple life cycle portfolio choice model.

1.2 The Difficulty in Separating Preferences and Frictions

1.2.1 Simple Life Cycle Model

We consider a standard life cycle portfolio choice model in the spirit of Cocco et al. (2005), which we describe in detail in Appendix A. investors in the model have time-separable preferences with discount factor β and a CRRA preferences over flow consumption with risk aversion σ. Each period investors earn exogenous stochastic labor income while they are working and retirement benefits while retired. In addition to making consumption and savings decisions, investors make choose the fraction of their wealth invested in a risky stock, θ_t, with the remaining fraction of their wealth, $1-\theta_t$, going into a risk-free bond. We assume a risk-free rate of 2% and a equity risk premium of 4.5%; the rest of the parameters are given in Appendix A.

We introduce two frictions into the model that affect investors portfolio choice decisions. First, there is a per-period participation cost p, which is incurred as a utility cost if $\theta_t > 0$ (e.g. Vissing-Jørgensen 2002; Fagereng et al. 2017; Catherine 2020; Briggs et al. 2021). This cost is designed to capture the costs associated with maintaining an account to invest in the risky asset, in addition to any hassle costs. Secondly, investors must pay a one-time portfolio to adjust their portfolio (e.g. Haliassos and Michaelides 2003; Gomes and Michaelides 2005; Abel et al. 2013), which is designed to capture the physical and opportunity costs associated with altering a portfolio allocation. Specifically, if investors choose $\theta_t \neq \theta_{d,t}$, where

$$\theta_{d,t} = \theta_{t-1} * \frac{R_t}{(1-\theta_{t-1})R_f + \theta_{t-1}R_t}, \tag{1}$$

they are required to pay a utility cost f. This expression captures that default asset allocations are equal to the asset allocation from the prior period, after adjusting for return realizations. We also assume $\theta_{d,0} = 0$ to capture the fact that the default for most investors when they begin working is non-participation in the stock market.

For different values of σ, f, and p, we solve the model numerically using standard techniques (Carroll 2020) to determine investors policy functions $\theta^*(\cdot)$, which depend on four state variables: age, wealth, income, and $\theta_{d,t}$. We then simulate the choices of investors and calculate average stock market participation over the life cycle.

1.2.2 Cross-Sectional Moments Don't Separate Preferences and Frictions

In Figure 2, we plot the life-cycle of participation for four different parameterizations of the model: (i) high σ and low f; (ii) low σ and high f; (iii) high σ and low p; (iv) low σ and high p. As illustrated in the figure, these lines are essentially on top of each other. Figure 2 also plots the life-cycle of participation from the 1989-2019 Survey of Consumer Finances, estimated using the approach in Deaton and Paxson (1994). Comparing the simulated life cycle of participation from the model with the that from the SCF, we find models with coefficient of relative risk aversion of 2.5 and 50 fit the data equivalently well[4], depending on how large participation or adjustment costs are.[5]

The results in Figure 2 illustrate that cross-sectional moments, such as the life-cycle of participation, cannot separately identify preferences and frictions. Of course we have not incorporated all possible forms of preferences into this model, but we view a model with a high coefficient of relative risk-aversion and a small friction as a reduced-form representation of other models. As discussed above, preference-based explanations for stock market non-participation consist of either introducing first-order risk-aversion in preferences, additional risks, or changing beliefs, all of which can be mapped into higher risk-aversion.

The implication of the results in this section is that separately identifying preferences and frictions requires additional variation beyond standard cross-sectional moments. In Section 2, we use variation from changes in 401(k) default asset allocations. Through the lens model in this section, this corresponds to finding variation in $\theta_{d,t}$ that induces changes in θ^* independently of $u(\cdot)$, which we use to identify f (or p) separately from σ.

[4]We are assessing model fit entirely on the life-cycle of participation or stock shares because we are interested in preferences over risk. However, because investors have time-separable CRRA preferences, these models will have drastically different implications for wealth accumulation, as σ controls the EIS. We could address this concern by using Epstein-Zin preferences and holding fixed the EIS, but do not for simplicity. We plan to incorporate Epstein-Zin preferences into model in Section 4 in future drafts.

[5]See Figure A1 for an analogous graph to Figure 2 for portfolio shares.

2 Data and Quasi-Experimental Variation

Section 1 shows cross-sectional moments can't separately identify investors' risk preferences from choice frictions. In this section, we describe the data and quasi-experimental variation we use to separately identify these two, which do using a theoretical framework in Section 3.

2.1 The Ideal Experiment

The ideal experiment for identifying preferences in the presence of frictions would be to randomly give some investors an investment account with risky assets. By randomly assigning accounts to investors, we would effectively remove any one-time participation (or adjustment) costs. Ideally, we would also remove any per-period costs associated with participation in the stock market, such as paying for the investment account, possibly by paying them for the investor.

In this ideal experiment, there are two potential outcomes with respect to stock market participation. First, if investors sell the risky assets, this would suggest they prefer safe assets over risky assets (e.g. due to loss aversion). Second, if investors kept the risky assets, this would suggest a preference for risky assets that is obscured by one-time or per-period participation costs. With respect to stock shares, there are more potential outcomes. However, by observing the fraction of risky assets investors sell, we would be able to infer their preferred stock share.

2.2 401(k) Administrative Data

The data we use to approximate this ideal experiment are 401(k) administrative records from a large U.S. financial institution, which is also used by Choukhmane (2021). These records cover around 600 employers (i.e. firms) and contain investor-level panel data from December 2006 to December 2017. For each employee in each year, we observe demographic characteristics, participation status in a 401(k) plan, 401(k) balances in dollars, employee and employer contributions in dollars, employee asset allocations by CUSIP, employer default asset allocations (also by CUSIP) recorded at the end of each month. When describing employees in our sample, we use the terms employee and investor simultaneously. In Table 1, we provide summary statistics on our data compared to those from the Survey of Consumer Finances (SCF) 2007, 2010, 2013, and 2016 waves. Relative to the median investor in the SCF, the median investor in our sample earns slightly less

income[6], is younger, and has slightly more retirement wealth. The latter is not surprising given that some households in the SCF don't have a retirement account and hence have a retirement wealth of zero.

A inherent limitation of our data is that we only observe the portfolio choices of investors in our sample within their 401(k) provided by their employer. We do not view this as a major limitation for three reasons. First, among all households in the SCF, only 5% participate in the stock market entirely outside their retirement account. Secondly, among households in the SCF who are eligible for a retirement account[7], 95% of the total equity holdings of the mean investor occurs just within their retirement accounts. Finally, the equity share of the average investor in the SCF is basically the same in their retirement account as in their total financial wealth, regardless of whether the investor is eligible for a retirement account. Taken together, these three findings suggest the equity holdings observed in our sample of retirement accounts are likely to provide a good picture of these investors' total equity holdings.

2.3 Description of Quasi-Experiments

We use two quasi-experiments motivated by the ideal experiment described above. In the first quasi-experiment, we compare the portfolio choices of employees hired within 12 months before and 12 months after 6 firms change the default asset allocation in their auto-enrollment 401(k) plans. The control group is 1,086 employees hired before the change, who are defaulted into a money market fund (i.e. no stock market exposure), and the treatment group are 1,321 investors hired *at the same firms* after the change, who are defaulted into a TDF (i.e. has stock market exposure). We call this the "money market to TDF" sample.[8]

Under the assumption that investors hired before and after the change are identical (and other assumptions formalized in the next subsection), this quasi-experiment provides a close approximation to the ideal experiment. In particular, the treatment group has zero one-time costs associated with stock market participation. Moreover, an advantage of this 401(k) setting is that there are

[6]For each investor in our sample, we observe their chosen contribution rate and their contribution amount. We use these two values to recover investors' income. However, this implies that we can only calculate median rather than mean income, since the income of investors not participating in 401(k) plans is not observed.

[7]We define SCF investors as being eligible for a retirement account if they report having access to a retirement account and/or they report assets in one. This is an imperfect measure for two reasons. First, some households may not realize they have access to a retirement account. Second, retirement accounts in the SCF includes accounts of all types beyond just 401(k) plans, such as 403(b) plans, IRAs, and other thrift plans.

[8]All six of these firms change their default asset allocation in 2007 following the passage of the Pension Protection Act of 2006.

no explicit per-period costs associated with maintaining or managing the account, in contrast to a brokerage account.[9] Therefore, if we observe the treatment group quickly re-balancing out of the TDF into the money market fund, this is likely to be driven by preferences rather than frictions.

In our second quasi-experiment, we compare the portfolio choices of employees hired within 12 months before and after 191 firms change their 401(k) plans from an opt-in regime to automatic-enrollment with a TDF as the default asset allocation. The control group is 40,337 employees hired before the change under the opt-in regime, while the treatment group is 52,400 employees hired after the change and automatically-enrolled into a TDF. Figure A2 displays the percent of the total number of firms that change their default in each year, which illustrates this variation is relatively evenly distributed between 2006 and 2017. We call this the "opt-in to TDF" sample.

Relative to the money market to TDF sample, the opt-in to TDF sample has the advantage of a much larger sample of firms and employees. However, an important difference is that in the opt-in to TDF sample is that the treatment and control groups differ both in terms of the frictions they face in investing in risky assets *and* the frictions they face in terms of saving in a 401(k) account. In principle, this is not a problem for identifying preferences for stock market participation, but simply means the frictions the control group must overcome to invest in risky assets might be larger.[10]

In Panels A and B of Figure 3, we plot the variation we use in both experiments. For the money market to TDF sample, Panel A plots 401(k) participation, money market participation, and stock market participation in the year when employees are hired based on the month in which investors are hired. investors in the left-half of each graph are in the control group, while investors in the right-half are in the treatment group. Consistent with a large literature on default effects (e.g. Madrian and Shea 2001; Chetty, Friedman, Leth-Petersen, Nielsen, and Olsen 2014; Blumenstock, Callen, and Ghani 2018; Beshears, Choi, Laibson, Madrian, and Skimmyhorn 2018), we find participation in the money market fund decreases discontinuously while stock market participation increases discontinuously for workers in the treatment group (i.e. hired right after the change in the default). In contrast, 401(k) participation remains unchanged. Panel B shows the analogous plot for the opt-in to TDF sample, in which we observe a discontinuous increase in 401(k) participation and stock market participation for investors in the treatment group.[11]

[9] There could of course still be per-period psychological costs, such as the cost of paying attention to managing the portfolio.

[10] Consistent with low per-period participation costs, we observe around 25,000 investors in the opt-in to TDF sample that have a positive stock share and account balance less than $500.

[11] In Figure A3, we compare investors in the treatment and control group in terms of age and income (the two other characteristics we can observe). In both samples we find the age of both groups are similar, while the income of the treatment group is similar in the money market to TDF sample and slightly higher in the opt-in to TDF sample.

2.4 Results from Quasi-Experiments

Figure 4 plots the results from our two quasi-experiments in separate panels.[12] In each panel, we plot the fraction of investors participating in the stock market and their average stock share τ years after they have been hired, where $\tau = 0$ corresponds to their choice immediately upon being hired. In both samples, we find almost all of the investors in the treatment group ($\approx 95\%$) continue holding some stock market exposure after they are hired. In contrast, investors in the control group gradually rebalance into stocks. In both samples, we also observe that investors in the treatment group maintain a relatively high stock share of around 80%, while the investors slowly increase their stock share as they rebalance into stocks.[13] Table 2 shows the economically meaningful difference in portfolio choices between the control and treatment groups in both experiments is also statistically significant and robust to the inclusion of firm fixed effects.

In sum, these results suggest most investors like to participate in the stock market, given the treatment group maintains their stock market exposure while the control group re-balances into stocks.[14] The difference in participation between the treatment and the control group indicates the existence of one-time adjustment costs associated with participating in the stock market, which could be real or psychological.

In order to make precise statements about investors' preferences, we need to place more structure on the responses to different default asset allocations in Figure 4. In the next section, we apply a theoretical framework developed by Goldin and Reck (2020) that allows us to convert results in Figure 4 into estimates of investors' preferences, taking into account the fact that frictions may affect investors' choices such that revealed preference fails.

[12] For the rest of the analysis, we only focus on portfolio choices made within 10 years of being hired. We drop choices after that since there are few investors who remain at the firm for that long.

[13] When calculating stock shares, we include both U.S. and international stocks. We use the portfolio allocations of the funds provided by the financial institution from which we obtained the data to convert holdings of mutual funds that include mixed allocations into stock share.

[14] We conduct several robustness checks on these results in Appendix C. Panel A of Figure A5, we show the responses of investors in the control group are similar to investors who we're hired in the years before them, suggesting the results are not driven by peer effects. In Panel B of Figure A5, we show the responses of investors in the control group are similar regardless of when the investor leaves the firm. This shows the increased stock market of the control group over time is not driven by a change in composition of preferences within the control group.

3 Identifying Risk Preferences Using 401(k) Default Switches

In this section, we describe and apply a framework developed by Goldin and Reck (2020) to identify the preferences investors in our sample over stock market participation and stock shares.

3.1 Theoretical Framework

Consider a continuum of investors indexed by i are hired by an employer at time $t = 0$ and make asset allocation choices at $t = 0, ..., T$. Denote $Y_{it} \in \{0, 1\}$ and $\theta_{it} \in [0, 1]$ as investor i's participation and stock share at time t respectively, where $Y_{it} = 1$ corresponds to participating in the stock market. We refer to t as investors' tenure, since it captures the length of time since the investor was hired. Each investor's participation and stock share decisions are subject to a time-invariant frame denoted by $D_i \in \{0, 1\}$, where $D_i = 1$ corresponds to an investor working for an employer with an auto-enrollment and a TDF as the default asset allocation (i.e. the treatment groups in both quasi-experiments) and $D_i = 0$ otherwise (i.e. the control groups). Throughout, we refer to D_i as the frame or default interchangeably. We also denote $\theta_i^d(D_i)$ as the default asset allocation faced by investor i, given frame D_i.

Each investor's preferred options at each tenure are denoted by $Y_{it}^* \in \{0, 1\}$ and $\theta_{it}^* \in [0, 1]$, which is not observed, while choices, denoted by Y_{it} and θ_{it}, are observed. Investors are characterized by a set of potential outcomes, $\{Y_{it}(d), \theta_{it}(d)\}_{d \in \{0,1\}}$, which generate their observe choices according to:[15]

$$Y_{it} = Y_{it}(d), \theta_{it} = \theta_{it}(d) \text{ if } D_i = d.$$

The primitives of this environment, $\left\{Y_{it}(0), Y_{it}(1), \theta_{it}(0), \theta_{it}(1), Y_{it}^*, \theta_{it}^*, D_i\right\}$, are assumed to be drawn from an identical population distribution with unrestricted dependence across i and t. As econometricians, we observe a panel of $(Y_{it}, \theta_{it}, D_i, age_{it})$, where $age_{it} \in \mathcal{A}$ denote an age group.

If an investor's participation or stock share decision is independent of the frame, we call that investor *consistent* with respect to that decision. Formally, we denote consistency by C_{it}^Y and C_{it}^θ,

[15] By writing choices as a function of potential outcomes, we are implicitly making a stable unit treatment value assumption (e.g. Rubin 1978) that investor i is not affected by the treatments of investors $j \neq i$. This is supported by the evidence in Panel A of Figure A5.

where

$$C_{it}^{Y} = \begin{cases} 1 & \text{if } Y_{it}(0) = Y_{it}(1), \\ 0 & \text{else.} \end{cases}$$

$$C_{it}^{\theta} = \begin{cases} 1 & \text{if } \theta_{it}(0) = \theta_{it}(1), \\ 0 & \text{else.} \end{cases}$$

In this framework, there are thus two possible types of investors for each decision: (i) consistent investors, whose choices are unaffected by frictions associated with the default; (ii) *in*consistent investors, whose preferences *are* affected by frictions associated with the default.

In what follows, we focus on identifying average preferences in the population at different tenures. Denote $E(\cdot)$ and $E_\tau(\cdot) = E(\cdot \mid t = \tau)$ as the expectation across the entire population and the conditional expectation across all investor at tenure $t = \tau$, respectively.

3.2 Bounding Average Preferences

3.2.1 Stock Market Participation

We begin by making the following assumptions, as in Goldin and Reck (2020).

Assumption 1 (Frame Separability)**.** *For all i and t, $(Y_{it}^*, \theta_{it}^*)$ is independent of D_i.*

Assumption 2 (Frame Exogeneity)**.** *D_i is independent of $(Y_{it}(0), Y_{it}(1), \theta_{it}(0), \theta_{it}(1))$.*

Assumption 3 (Frame Monotonicity)**.** *For all i and t,*

$$Y_{it}(1) \geq Y_{it}(0), \quad \theta_{it}(1) \geq \theta_{it}(0).$$

Assumption 4 (Consistency Principle)**.** *For all i and t,*

$$C_{it}^{Y} = 1 \implies Y_{it} = Y_{it}^*, \quad C_{it}^{\theta} = 1 \implies \theta_{it} = \theta_{it}^*$$

Intuitively, frame separability, requires investors' preferences to be independent of whether which default (or frame) they are given. This assumption results out the possibility that investors view the default as providing information and prefer to participate more when the default is partic-

ipation.[16] Frame exogeneity requires the choice of default to be independent of investors' preferences, or equivalently that investors in the treatment and control group have the same preferences. This is a natural assumption in our setting, given the default is chosen by the firm. In the case of participation, frame monotonicity requires that there are no investors who always choose the opposite of the default, which we view as reasonable given the literature on default effects. Finally, the consistency principle requires that consistent investors reveal their preferences. For example, if an investor choose to participate regardless of the default asset allocation, we assume the investor's preference is indeed to participate.

Under the previous assumptions, the following partial identification result from Goldin and Reck (2020) allows us to bound the preferences for stock market participation of the population of investors in our sample.

Proposition 1. *Under Assumptions 1-4, the average population preference for stock market participation among investors with tenure $t = \tau$ is partially-identified:*

$$E_\tau\left(Y_{it}^*\right) \in \left[E_\tau\left(Y_{it} \mid D_i = 0\right), E_\tau\left(Y_{it} \mid D_i = 1\right)\right]. \tag{2}$$

All derivations are provided in Appendix B. The intuition for this result is as follows. The average preference for stock market participation is the weighted-average of the preferences of consistent and inconsistent investors. Since Assumption 4 implies the preferences of consistent investors are revealed by their choices, we just need to bound the preferences of inconsistent investors to bound the population average preference. The bounds in (2) comes from considering the "worst-case" scenario, in which inconsistent investors preferences align with their defaults, in which case the bound is (2).

According to Proposition 1, the results in the left panel of Figure 4 provide the required information to bound average preferences for stock market participation. By (2), average preferences for stock market participation among all investors in our population lie somewhere between the choices of the treatment and control groups, which is illustrated in the left panel of Figure 5 for the money market to TDF sample.[17] For example, after investors have been at the firm for 3 years ($\tau = 3$), we can bound the fraction of investors who prefer stock market participation between 78% and 95%. This bound, which is strictly higher than average participation in our sample of 68% (Table 1), illustrates that one-time adjustment costs are driving a wedge between observed choices and

[16] Informally, the partial-identification results we present below are robust to allowing the frame to affect preferences monotonically. However, once preferences depend on frames, identification of preferences is no longer a well-defined problem because investors do not have an underlying stable preference relation.

[17] As evident from Figure 4, we find similar results across both samples.

underlying preferences. As tenure increases, more investors in the control group become consistent and reveal their preferences, resulting in a tighter bound.

3.2.2 Stock Shares

We now turn to identification of investors' preferences for stock shares. Unlike the previous section, Assumptions 1-4 are not sufficient to place meaningful bounds on the average preferences because stock shares are continuous variables. To see why this poses a problem, consider a hypothetical investor with $0 < \theta_{i\tau}(0) < \theta_{i\tau}(1) < 1$, for some $\tau \geq 0$. This investor is by definition inconsistent at τ, implying we cannot infer anything about his preferences from his observed choices under Assumptions 1-4. For example, this investor may have $\theta^*_{i\tau} \in (0, \theta_{i\tau}(0))$. The possible presence of such an investor renders the analogous bound to Proposition 1 for stock shares invalid, since we would have $E_\tau(\theta^*_{it}) < E_\tau(\theta_{it} \mid D_i = 0)$ in a world with investors of this type.

The prior counter-example illustrates the need to place more structure on the choices of inconsistent investors because investors can be inconsistent in an infinite number of ways with a continuous choice (unlike in the case of participation, which is a binary decision). Following an extension developed in the Online Appendix to Goldin and Reck (2020), we introduce the notion of partition-consistency. In particular, we define an investor at time t as *partition-consistent at* x if $C^\theta_{it}(x) = 1$, where

$$C^\theta_{it}(x) = \begin{cases} 1 & \text{if } \theta_{it}(0) \leq x \text{ and } \theta_{it}(1) \leq x, \\ 1 & \text{if } \theta_{it}(0) > x \text{ and } \theta_{it}(1) > x, \\ 0 & \text{else.} \end{cases}$$

Intuitively, an investor is partition consistent at x if they would behave consistently in a hypothetical choice problem between two intervals: $[0, x]$ and $(x, 1]$. We then make the following assumption, which is a generalization of Assumption 4.

Assumption 4′ (Partition-Consistency Principle)**.** *For all i, t, and x,* $C^\theta_{it}(x) = 1$ *implies*

$$\theta_{it} \geq x \implies \theta^*_{it} \geq x,$$
$$\theta_{it} \leq x \implies \theta^*_{it} \leq x.$$

This partition consistency principle is an intuitively states that investors who would choose a stock share above x regardless of D_i indeed prefers a stock share above x. Importantly, this allows an

investor who deviates from a default asset allocation to not necessarily reveal her true preferences (e.g. as in models with convex adjustment costs). Under this assumption, we can bound population preferences for stock shares analogously to the stock market participation bounds above using the following result.

Proposition 2. *Under Assumptions 1-3 and 4′, the average population preferred stock share among investors with tenure $t = \tau$ is partially-identified:*

$$E_\tau(\theta_{it}^*) \in [E_\tau(\theta_{it} \mid D_i = 0), E_\tau(\theta_{it} \mid D_i = 1)]. \tag{3}$$

Proposition 2 shows the results in the right panel of Figure 4 allow us to bound average preferences for stock shares, as illustrated in the right panel of Figure 5. After investors have been at the firm for three years, we can bound the average preferred stock share among all investors in our sample between 59% and 81% (in the money market to TDF sample). Note that this average preferred stock share is higher than the average stock share in the SCF of around 23% (Table 1) and the average stock share in our sample of 53%. Thus, these results again illustrate how frictions can drive a wedge between choices and preferences. Economically, in a life cycle portfolio choice model with CRRA preferences, no labor income risk, and a constant investment opportunity set (e.g. Merton 1969), these bounds imply bounds on the coefficient of relative risk aversion between 2.65 and 3.65.[18] and Compared to estimates in Figure 1, this is quite low, providing a preliminary illustration of how taking into account choice frictions can change inferences about risk preferences.

3.3 Estimating the Preferences of Consistent Investors

The previous section shows how we can bound the average preferences in the population using our quasi-experimental variation. These average preferences are the weighted-average of the preferences of consistent and inconsistent investors. In this section, we recover point-estimates of consistent investors' preferences. Recovering these preferences for stock shares requires the following additional assumption.

Assumption 5. *For all i, t, and x, $C_{it}^\theta(x) = 1 \implies C_{it}^\theta = 1$.*

This assumption states that all investors that are partition-consistent when choosing their stock shares are also consistent. Economically, this assumption is consistent with a large class of models

[18] In this calculation we assume annual risk premium of 5.5% and an annualized standard deviation of 16%.

of default effects in which investors' preferences can be represented as-if deviating from a default requires a one-time adjustment cost (see Masatlioglu and Ok 2005, for an axiomatization). However, this assumption does rule out some models, such as those with convex adjustment costs or where default effects are driven by cognitive uncertainty (e.g. Gabaix 2019; Enke and Graeber 2020). Under this assumption, the following proposition shows how we can recover the preferences of consistent individuals.

Proposition 3. *Under Assumptions 1-4,*

$$E_\tau(Y_{it}^* \mid C_{it}^Y = 1) = E_\tau(Y_{it} \mid Y_{it} \neq D_i). \tag{4}$$

Under Assumptions 1-3, 4′, and 5,

$$E_\tau(\theta_{it}^* \mid C_{it}^\theta = 1) = E_\tau(\theta_{it} \mid \theta_{it} \neq \theta_i^d(D_i)). \tag{5}$$

Proposition 3 states that we can recover an estimate of the average preference for participation and stock shares among consistent investors by simply looking at the choices of investors who deviate from the defaults. Intuitively, we can do this because the consistency principle implies consistent investors choices reveal their preferences. Although not *all* consistent investors deviate from the default, the fact that D_i is set by the firm and uncorrelated with investors' preferences (by frame exogeneity) ensures the preferences of those who do deviate are similar to those who don't.

Figure 6 plots our estimate of the preferences of consistent investors. In both samples, after three years of tenure, we find about 94% of consistent investors prefer stock market participation with an average stock share around 75%.[19] Consistent with our population bounds, these both are quite high, suggesting these consistent investors indeed like stocks but face one-time frictions associated with participation. The fact that our estimates are similar across the two samples provides support for the consistency principle: even though consistent investors face frictions of different sizes in the two samples, when overcoming they reveal similar preferences.

We can also use Proposition 3 to estimate the fraction of consistent investors by looking at the fraction of investors who deviate from the default. Figure A6 plots how this varies with age. We find the fraction of consistent investors is slightly increasing over the life cycle, which is consistent with a one-time real adjustment cost which younger investors (who have less wealth) are less likely to pay.[20] Figure A6 also shows that a 40 year-old investor (the average in our sample) is

[19] Figure A8 plots estimates of these preferences based on what default consistent investors were given. For stock shares, we find consistent investors reveal similar preferences regardless of the default, which supports Assumption 5.

[20] The fact that consistency increases over the life cycle is also consistent with theories of financial literacy accumu-

more likely to be consistent in the money market to TDF sample than in the opt-in to TDF sample. This is consistent with the fact that the frictions the control group needs to overcome in the latter sample are likely large, due to a need to change both the default asset allocation and opt-in to a 401(k) account to begin with. In Figure A7, we plot how the fraction of consistent investors varies by default: investors are less likely to be consistent when they are defaulted into a TDF.

3.4 Estimating Average Preferences

In this section, we estimate the average preferences among investors in the population. Without any assumptions, we can derive expressions for the relationship between the preferences consistent investors and the preferences of the population:

$$E_\tau(Y_{it}^*) = E_\tau\left(Y_{it}^* \mid C_{it}^Y = 1\right) - \frac{1}{E_\tau(C_{it}^Y)} \underbrace{cov_\tau\left(Y_{it}^*, C_{it}^Y\right)}_{\text{selection bias}}, \tag{6}$$

$$E_\tau(\theta_{it}^*) = E_\tau\left(\theta_{it}^* \mid C_{it}^\theta = 1\right) - \frac{1}{E_\tau(C_{it}^\theta)} \underbrace{cov_\tau\left(\theta_{it}^*, C_{it}^\theta\right)}_{\text{selection bias}}. \tag{7}$$

Proposition 3 shows the first two terms in (6) and (7) are identified, which we estimated in Figure 6. However, the second term in both expressions is not identified. Intuitively, it represents a form of selection bias that arises if consistent investors have different preferences from inconsistent investors. In general, this selection bias is unbounded without placing further restrictions on investor decision-making. For example, in our structural model described in Section 4, we could compute this directly.

To point-estimate average population preferences, we begin by making the following (strong) identifying assumption.

Assumption 6. *For all i and τ,*

$$cov_\tau(Y_{it}^*, C_{it}^Y) = cov_\tau(\theta_{it}^*, C_{it}^\theta) = 0.$$

This assumption states that once we condition on an investor's tenure, whether they are consistent is uncorrelated with their preferences. We view this as a reasonable starting assumption because income risk is a key driver of portfolio choice in standard life cycle models, while empirical ev-

lation (e.g. Lusardi, Michaud, and Mitchell 2017).

idence suggests income risk is highly dependent on tenure (e.g. Kogan, Papanikolaou, Schmidt, and Song 2020). However, even if this assumption is unreasonable, we can still rely on our bounds of population preferences given in Figure 4.

Under Assumption 6, the preferences of the population at each tenure are given by the preferences of consistent investors in Figure 6: at tenure $\tau = 3$, the average preference for participation is 94% and the average preferred stock share is 75%. Using the back-of-the-envelope calibration from above with the Merton (1969) model, this latter suggests an estimate of relative risk aversion of 2.8. Consistent with our bounds above, this suggests investors in our sample prefer stock market participation with reasonably high stock shares, but face one-time participation costs associated with doing so.

3.5 Estimating Preferences over the Life Cycle

Assumption 6 is relatively strong in that it rules out the possibility that consistency and preferences might both vary with age (conditional on tenure). This is restrictive given the the stock of human capital, the central driver of portfolio choice in most life cycle models, decreases with age, while there are natural reasons to believe consistency might vary with age as well (e.g. older investors have a lower opportunity cost). We thus relax Assumption 6 by making the following assumption.

Assumption 7. *For all i, τ, and $A \in \mathcal{A}$,*

$$cov_\tau(Y_{it}^*, C_{it}^Y \mid age_{it} = A) = cov_\tau(\theta_{it}^*, C_{it}^\theta \mid age_{it} = A) = 0.$$

Assumption 7 is a weaker version of Assumption 6 in that it conditions on age in addition to tenure. Under Assumption 7, we can identify how preferences vary over the life cycle. In particular, we can apply the law of iterated expectations to (6) and (7) to obtain:

$$\begin{aligned} E_\tau(Y_{it}^* \mid age_{it} = A) &= E_\tau\left(Y_{it}^* \mid C_{it}^Y = 1, age_{it} = A\right), \\ E_\tau(\theta_{it}^* \mid age_{it} = A) &= E_\tau\left(\theta_{it}^* \mid C_{it}^\theta = 1, age_{it} = A\right). \end{aligned}$$

Applying Proposition 3 conditionally then gives the following life cycles of preferences for investors of tenure $t = \tau$:

$$E_\tau(Y_{it}^* \mid age_{it} = A) = E_\tau\left(Y_{it} \mid Y_{it} \neq D_i, age_{it} = A\right), \tag{8}$$

$$E_\tau(\theta^*_{it} \mid age_{it} = A) = E_\tau\left(\theta_{it} \mid \theta_{it} \neq \theta^d_i(D_i), age_{it} = A\right). \tag{9}$$

Panel A of Figure 7 plots our estimates of participation preferences over the life cycle using (8) by tenure for the opt-in to TDF sample. Regardless of what tenure we use, we find investors have similar preferences for participation that are high and relatively flat over the life cycle. We find some evidence of participation increasing between ages 20 and 30, but it remains flat above 90% after that. For comparison, Panel B of Figure 7 plots the observed participation in our sample in 2007 and in the SCF 2007 wave. In 2007, the life cycle of *observed* participation is increasing and strictly below 70%, similar to the SCF.

Figure 8 shows the analogous results for stock shares. Across all tenures, Panel A shows the average preferred stock share among investors in our sample is relatively high (strictly above 60%) and mostly decreasing over the life cycle. Notably, this life cycle of preferences is consistent with standard predictions from life cycle portfolio choice models with risky labor income that is uncorrelated with stock returns (Campbell and Viceira 2001). As in the case of participation, we find evidence that investors' preferred stock share increases slightly in the first between ages 20 and 30. Moreover, investors' preferred stock share over the life cycle looks substantially different than their observed stock shares shown in Panel B, which is lower and hump-shaped with age (as in Ameriks and Zeldes 2004).

Figure 9 plots our estimates of investors' average preferences over the life cycle across all tenures compared to the glide path of the TDF in our sample. The results show investors prefer a lower stock share than the allocation of a TDF. When investors are below age 35, this difference is particularly stark as investors' preferred stock share is upward sloping, while the life cycle of a TDF is flat before 35. This preference of younger investors for a lower stock share is also evident from Figure A13: consistent investors defaulted into a TDF choose a lower stock shares at all ages, but this is especially true for younger investors.

Collectively, these results suggest that in the absence of participation and adjustment costs, investors in our sample prefer to participate in the stock market with reasonably high stock shares. The stark difference between observed choices and estimated preferences in Figures 7 and 8 highlights the importance of distinguishing between the two in the presence of choice frictions. These findings have implications for the large literature studying limited household stock market participation. If investors in our sample had preferences with first-order risk-aversion, had pessimistic beliefs, or faced additional risks associated with participation, our estimation would have recovered a low preference for participation. The fact that, after taking into account choice frictions, we estimate investors prefer participate in the stock market provides support for explanations of low

observed participation based on real or psychological participation and adjustment costs.

3.6 Robustness

Preference heterogeneity. In Figure A9, we explore preference heterogeneity[21] over the life cycle by plotting the distribution of preferred stock shares by consistent investors for three different age groups: 20-34, 35-49, and over 50. These three groups are roughly evenly-spaced terciles. We find preference heterogeneity increases over the life-cycle: most investors in the lowest age group prefer a stock share of over 80%, while there is much more dispersion in preferred stock shares among the highest age group.

Year, cohort, and firm effects. In Figures A10 and A11 we test for the presence of year and cohort effects in investors' choices within our sample. We find evidence that the life cycles of participation and stock shares has increased slightly over time, especially for investors hired under auto-enrollment. This is consistent with the shift away from money market funds to TDF as the default asset allocation in 401(k) plans following the Pension Protection Act of 2006 (e.g. Campbell 2016). We also find some evidence of an increase in participation and stock shares in younger cohorts. Together, these results indicate the sum of the changes in cohort and time effects is positive, which suggest our estimates might under-estimate the preferences of older investors and over-estimate the preferences of younger investors. In Figure A12, we show our estimates of preferences are robust to including firm fixed effects.

Conditioning on income. In our data, we observe the salary investors receive from their employer for the subset of investors who contribute to their 401(k) plans. Thus, we can estimate average preferences under a weaker version of Assumption 7, where we assume that consistency and preferences are uncorrelated conditional on age, tenure, *and* income. In Figure A14, we plot estimates of preferences over the life cycle at different tenures using this weaker assumption. The results show that our estimates of preferences are unaffected. In Figure A15, we plot our estimates of preferences over the life cycle by income quartiles, after integrating over tenure. The results show our estimates our preferences are mostly similar across income quartiles, consistent with the results in Figure A14. However, we find some evidence that investors in the bottom income quartile like stocks slightly less than those in the top three quartiles. This could be driven by many explanations, but it is consistent with other evidence that suggests non-homotheticity in preferences (Brunnermeier and Nagel 2008; Straub 2019; Meeuwis 2020).

[21] We use the term preference heterogeneity loosely, as these could also reflect heterogeneity in beliefs as in Giglio, Maggiori, Stroebel, and Utkus (2021).

4 Life Cycle Portfolio Choice Model

In this section, we describe a rich life cycle portfolio choice model and estimate it using the variation from our quasi-experiments. In this model, investors choose a level of consumption, retirement wealth, liquid wealth, unsecured debt, and the portfolio allocation in their retirement wealth. This model, which builds on Choukhmane (2021), has three key features that are required to match our quasi-experimental evidence. First, investors can choose different asset allocations for existing wealth and new contributions in their retirement account. Second, investors are required to pay separate opt-out costs to deviate from the default contribution rate in their retirement account and the default portfolio allocation. When agents are hired, these defaults are employer-specified; in later periods, choices from the prior period are the current period's default. Finally, investors face uncertainty about future earnings and employment status at their current employer that creates value for investors to delaying portfolio adjustments and savings decisions from their defaults. In Appendix D, we summarize the model parameters.

4.1 Model Description

4.1.1 Demographics and Preferences

Investors are born at $t = 0$ and work T^w periods with their first retirement period at $t = T^w$. Each period corresponds to one year. Investors die with certainty at $t = \overline{T}$, at which point all of their resources are bequested, such that investor's last period in which they can consume is $t = \overline{T} - 1$. Before their certain death, investors survival probability is time-varying and denoted m_t, which is taken from the SSA. Denote investor's age as $a_t = t + a_0$, where a_0 is the age at which investors are born.

Investors have time-separable expected utility preferences with a CRRA Bernoulli utility function over consumption. Denote investors' annualized time discount factor as β and their coefficient of relative risk aversion (or equivalently inverse of elasticity of intertemporal substitution) as σ. Per-period flow utility is adjusted for an equivalence scale such that it is equal to

$$u_t(c) = n_t * \frac{\left(c/n_t\right)^{1-\sigma}}{1-\sigma},$$

where n_t is taken from Lusardi et al. (2017).

4.1.2 Labor Income

There is a continuum of employers indexed by e for which investors can work. At any point in time, an investor's employment status, denoted emp_t, is in one of four states: E = employed in the current and prior period; JJ = employed in a different job in the current period from prior period; U = unemployed in the current period; Ret = retired. When investors are born at $t = 0$, they are randomly assigned to an employer with $emp_0 = E$. When investors are employed, we also keep track of their tenure, denoted by ten_t.

The current state of employment determines the income process the investor's receive in the current period in addition to the transition probabilities across states in the next period. We now describe these different income processes in turn.

Employment: $emp_t = E$**.** While working, investors supply one unit of labor inelastically and earn an income w_t that is stochastic and exogenous. This income consists of a deterministic component that is cubic in age and a stochastic component that follows an AR(1) process with normal innovations:

$$\begin{aligned} \ln w_t =& \delta_0 + \delta_1 a_t + \delta_2 a_t^2 + \delta_3 a_t^3 + \eta_t, \quad \eta_t = \rho \eta_{t-1} + \xi_t^E, \\ & \xi_0^E \sim N(0, \sigma_{\xi_0}^2), \quad \xi_t^E \sim N(0, \sigma_\xi^2) \ \forall t > 0. \end{aligned} \tag{10}$$

Investors' tenure status evolves according to $ten_t = ten_{t-1} + 1$. Investors do not receive any labor income in the event that they die in the current period. Additionally, when investors are born at $t = 0$, the distribution of η_t^E is different to account for heterogeneity in the initial period income shock.

Job transition: $emp_t = JJ$**.** In each period, investors can switch jobs. We model these transitions separately because of the fact that retirements are employer-specific, so we need to account for when investors change jobs. If they do so, their income evolves according to :

$$\ln w_t = \delta_0 + \delta_1 a_t + \delta_2 a_t^2 + \delta_3 a_t^3 + \eta_t, \quad \eta_t = \rho \eta_{t-1} + \xi_t^{JJ}, \quad \xi_t^{JJ} \sim N(\mu^{JJ}, \sigma_\xi^2). \tag{11}$$

This earnings process captures a wage premium associated with switching jobs. Investors' tenure is reset to $ten_t = 0$ during a job transition.

Unemployment: $emp_t = U$**.** When investors are unemployed, they receive unemployment benefits equal to $ui_t = ui(\eta_t)$, where $ui(\eta_t)$ is described below. When hired back after being unem-

ployed, next-period income evolves according to

$$\ln w_t = \delta_0 + \delta_1 a_t + \delta_2 a_t^2 + \delta_3 a_t^3 + \eta_t, \quad \eta_t = \rho \eta_{t-1} + \xi_t^U, \quad \xi_t^U \sim N(\mu^{UE}, \sigma_\xi^2). \tag{12}$$

This earnings process captures a wage reduction from switching jobs. Investors' tenure is irrelevant in this state.

Retirement: $emp_t = Ret$**.** When $t \in [T^w, \overline{T} - 1]$, investors are retired and earn retirement benefits denoted by ss_t, which are described below. Note that investors do not earn any retirement benefits at the time of death. Investors' tenure is irrelevant in this state.

State transitions. Investors' employment status, emp_t, evolves according to a first-order Markov chain. States are ordered as follows: E, JJ, U, Ret. The state transition matrix, which is tenure and age-dependent, is:[22]

$$\begin{pmatrix} 1-\pi^{JJ}(t,ten_t) & \pi^{JJ}(t,ten_t) & \pi^{EU}(t) & 0 \\ 1-\pi^{JJ}(t,ten_t) & \pi^{JJ}(t,ten_t) & \pi^{EU}(t) & 0 \\ 0 & \pi^{UE}(t) & 1-\pi^{UE}(t) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

4.1.3 Financial Assets

There are four financial assets in the model. First, there is a risk-free bond that has a constant gross return of $R_t^B = R_f$ per year. Second, there is a risky asset, a stock, that corresponding to a diversified market index and pays a stochastic IID gross return of $R_t^S = R_t$ per year, where

$$\ln R_t^S = \ln R_f + \mu_s + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_s^2). \tag{13}$$

Third, investors have access to a liquid risk-free asset that has a constant gross return of $1 + r$ per year. This corresponds to the interest paid on short-term checking and savings accounts. Finally, there is unsecured credit, which is subject to an interest rate of $1 + r^{CC}$ and subject to a borrowing limit.

[22] Entry i-j of this matrix corresponds to the probability of transitioning from the i-th state at t to the j-th state at $t+1$

4.1.4 Savings Accounts

The asset side of investors' balance sheets are comprised of two savings accounts, which we now describe in turn.

Liquid savings account. When investors are born at $t = 0$, they are endowed with zero liquid wealth and a liquid savings account. Investors can only use this liquid savings account to invest into a liquid riskless asset or borrow using unsecured credit. The borrowing limit is $\underline{L}_t$. Returns in this account are subject to capital taxation at rate τ^k. The balance of this account, denoted by L_t, evolves according to:

$$L_{t+1} = \left(L_t + s_t^l\right)\left[1 + r(L_t)\right], \quad L_0 = 0, \tag{14}$$

where s_t^l is the net savings the investor places in this account and

$$r(L) = \begin{cases} 1 + r\left(1 - \tau^k\right) & \text{if } L \geq 0, \\ 1 + r^{cc} & \text{else .} \end{cases}$$

Defined-contribution retirement savings account. When investors are born at $t = 0$, they are endowed with one savings account with their current employer to which they are initially matched. This savings account is the counterpart in our model to a defined-contribution 401(k) retirement savings plan. Investors use this savings account to buy and sell any combination of the bond and the stock, subject to the restriction of no margin trading (i.e. no leveraged purchases or short-sales). Investors cannot use this account to purchase the liquid risk-free asset or borrow. Returns earned in this account are tax free.

The balance of this savings account, denoted by A_t, evolves according to:

$$A_{t+1} = A_t \sum_{j \in \{B,S\}} \Theta_t^j R_{t+1}^j + \mathcal{M}_e(s_{t+1}^{dc}, t, ten_t, w_t), \quad A_0 = 0, \tag{15}$$

$$\begin{aligned} \mathcal{M}_e(s, t, ten_t, w_t, emp_t) = s * w_t + 1(emp_t \in \{E, JJ\}) \times w_t \Big(& fixed_e \\ & + \Upsilon_e(t, ten_t) \times match_e(ten_t) \times \min\{s, cap_e\} \Big) \end{aligned} \tag{16}$$

where s_t^{dc} is the net savings rate (including possible withdrawals) the investor places in this account, $\mathcal{M}_e(\cdot)$ is an employer-specific function that determines how savings is mapped into account contributions, and $\left\{\Theta_t^B, \Theta_t^S\right\} \in \mathbb{R}_+^2$ are the portfolio shares of existing assets in stocks and bonds respectively chosen at time t. The form of $\mathcal{M}_e(\cdot)$ captures the fact that when placing s into the retirement account, investors benefit from employer-specific matching. This employer match is characterized

by a level of non-contingent contributions ($fixed_e$), a tenure-dependent match rate ($match_e(\cdot)$), and a threshold contribution rate (cap_e). Additionally, we adjust these employer matches by a factor $\Upsilon_e(\cdot) \leq 1$ to capture the possible loss of employer matches if investors separate from the employer before vesting is complete.

In addition to choosing $\left\{\Theta_t^B, \Theta_t^S\right\}$, investors also get to choose $\left\{\theta_t^B, \theta_t^S\right\} \in [0,1]^2$, which correspond to the share of new contributions invested asset j. Thus, the balance of investors' savings accounts depends on *two* portfolio choice decisions: (i) their portfolio allocation of existing contributions and (ii) their portfolio allocation of new contributions. This distinction, which matches the institutional features of 401k asset allocation decisions, is important in our model because investors are subject to default effects (described in Section 4.1.5). Without the presence of such default effects, these two decisions could be collapsed into one portfolio choice decision. For notational convenience, denote Ξ_t as the four-dimensional vector of portfolio choices for the two assets.

Defined-contribution account roll-overs. Because the evolution of the savings account is employer-specific ($\mathcal{M}_e$ depends on e), we need to specify what happens to 401k retirement savings plans when investors switch jobs or become unemployed. When investors become unemployed, we assume investors maintain access to this account, but cannot make any more contributions to the account, as represented by the constraint $s_t^{dc} \leq 0$. When investors switch jobs, we assume investors maintain access to their existing account, but moving forward the employer matching function, denoted by $\mathcal{M}_e(\cdot)$ in (15), and the default asset allocation for new contributions, described in Section 4.1.5, change to those specified by the new employer.

4.1.5 Default Effects

An investor's portfolio allocation and savings decisions in the defined-contribution account are both subject to default effects. We first describe the value of these defaults and then how they impact the investor's choices.

Default asset allocation for new DC contributions. When an investor begins working for employer e at time t, the default asset share of contributions to the defined-contribution savings account invested in asset j is $\overline{\theta}_e^j$. These employer-specific default asset allocations are drawn from an exogenous distribution of employer types, with probability mass function by $F_{\overline{\theta}}(\cdot)$. Later in the worker's tenure, the default asset allocation for contributions corresponds to the allocation chosen

in the prior period. Formally, for $j \in \{B, S\}$,

$$\theta_{d,t}^{j} = \begin{cases} \overline{\theta}_{e}^{j} \sim F_{\overline{\theta}}(\cdot) & \text{if } emp_t = JJ, \\ \theta_{t-1}^{j} & \text{else.} \end{cases} \tag{17}$$

Default portfolio allocation for existing DC contributions. When choosing the portfolio allocations of existing assets, the default allocation for each asset is equal to the amount of old contributions in that asset, adjusted for realized returns, plus the amount of new contributions allocated to that asset. Formally, for $j \in \{B, S\}$,

$$\Theta_{d,t}^{j} = \begin{cases} \frac{A_{t-1}\Theta_{t-1}^{j}R_t^{j} + \mathcal{M}_e(s_{t-1}^{dc})\theta_{t-1}^{j}R_t^{j}}{A_{t-1}\sum_j \Theta_{t-1}^{j}R_t^{j} + \mathcal{M}_e(s_{t-1}^{dc})\sum_j \theta_{t-1}^{j}R_t^{j}} & \text{if } s_{t-1}^{dc} > 0, \\ \frac{A_{t-1}\Theta_{t-1}^{j}R_t^{j}}{A_{t-1}\sum_j \Theta_{t-1}^{j}R_t^{j}} & \text{else,} \end{cases} \tag{18}$$

Note that specification embeds the assumption that when investors dis-save out of their DC account, they sell assets in proportion to their current portfolio allocations. Additionally, we set $\Theta_{d,0}^{j} = 0$, since investors are born with no assets.

Default contribution rate in DC account. When an investor begins working for employer e at time t, the default contribution rate their defined-contribution savings account is $\overline{s}_e^{dc}$. The three employer-specific default asset allocations are drawn from an exogenous distribution of employer types, with probability mass function by $F_{\overline{s}}(\cdot)$. Later in the worker's tenure, the default contribution rate is equal to the contribution rate from the prior period. Formally,

$$s_{d,t} = \begin{cases} \overline{s}_e^{dc} \sim F_{\overline{s}}(\cdot) & \text{if } emp_t = JJ, \\ s_{t-1}^{dc} & \text{else.} \end{cases} \tag{19}$$

Effect of defaults. Investors in our model are subject to default effects when making portfolio choice decisions and savings decisions in the defined-contribution account. Recall Ξ_t denotes the vector of investor's portfolio allocations in the defined-contribution account. If an investor chooses $\Xi_t \neq \Xi_{d,t}$, where

$$\Xi_{d,t} = \left(\Theta_{d,t}^{B}, \Theta_{d,t}^{S}, \theta_{d,t}^{B}, \theta_{d,t}^{S}\right),$$

the investor incurs a utility cost k_θ. This cost is designed to capture any physical or physiological costs associated with making portfolio choice decisions. Similarly, choosing $s_t^{dc} \neq s_{d,t}$ requires incurring a utility cost of k_s.

4.1.6 Government

Income taxes. Investors face a non-linear income tax schedule $tax_i(w_t)$. Contributions to the DC retirement account are not subject to income taxation, while withdrawals (in either unemployment or retirement) increase taxable income by the withdrawal amount.[23]

Unemployment benefits. Investors receive an unemployment benefit of $ui(\eta_t)$ when uncompleted. This benefit depends on labor productivity, η_t, from the last period in which the agent was employed. Any withdrawals from the DC retirement account are treated as compensation and thus may offset unemployment benefits.

Retirement benefits. After retirement, investors receive benefits modeled after Social Security, denoted by $ss(ae_{T^w})$, where ae_{T^w} is the investor's average lifetime earnings at the time of retirement, which evolves according to:

$$ae_{t+1} = \begin{cases} \frac{w_{t+1}+a_t * ae_t}{a_t+1}, & \text{if } t < T^w, \\ ae_{T^w} & \text{else.} \end{cases}$$

4.1.7 Recursive Formulation

Investors face a dynamic optimization problem with 12 state variables: a_t = age; η_t = labor productivity, emp_t = employment status; e = employer; ten_t = tenure; ae_t = average lifetime income; A_t = DC retirement savings; L_t = liquid savings; $\Xi_{d,t} \in \mathbb{R}^4$ = default portfolio shares; $s_{d,t}$ = default contribution to DC account. Using the fact that portfolio shares sum to one, this can be reduced to a problem with 10 state variables by dropping the portfolio shares in the bond. Denote the vector of these state variables as X_t.

In this optimization problem, investors have 7 controls: c_t = consumption; $\Xi_t \in \mathbb{R}^4$ = portfolio shares; s_t^{dc} = defined-contribution savings rate; s_t^l = liquid savings. As above, this can be reduced to 4 controls given portfolio shares sum to one and consumption is pinned down by the budget constraint. In choosing these controls, investors are restricted from borrowing and engaging in any margin trading (i.e. no short-selling or leveraged positions):

$$A_t \geq 0, \quad L_t \geq \underline{L}_t, \quad \Theta_t^j \in [0,1], \quad \theta_t^j \in [0,1], \quad \sum_j \Theta_t^j = \sum_j \theta_t^j = 1. \tag{20}$$

[23] Thus, the DC account in our model is modeled after the traditional tax-deferred DC model rather than the Roth-401(k) model.

We now characterize the value function of an investor, $V(\cdot)$, separately for the four states of employment emp_t.

Retirement: $emp_t = Ret$**.** There are two sources of uncertainty when decisions are made at time t: mortality occurring with probability m_{t+1}; ε_{t+1} = asset return shocks. An investor's value function is thus characterized by the following recursive equation:

$$V(X_t) = \max_{s_t^{dc}, s_t^l, \Xi_t} \left\{ u_t\left(c_t - k_\theta * 1\left\{\Xi_t \neq \Xi_{d,t}\right\}\right) + (1 - m_{t+1})\beta E\left[V(X_{t+1}) \mid X_t\right] \right\},$$

subject to: (13), (14), (15), (17), (18), (20), and

$$s_t^l = ss_t - s_t^{dc} * w_t - c_t,$$
$$V(a_{\overline{T}}, \cdot) = 0,$$
$$s_t^{dc} \leq 0.$$

Working-life: $emp_t \in \{E, JJ\}$**.** There are five sources of uncertainty when decisions are made at time t: mortality occurring with probability = m_{t+1}; ε_{t+1} = asset return shocks; employment risk based on the state transition matrix; labor income shocks based on ξ_{t+1}^E or ξ_{t+1}^{JJ}; the type of future employer after a job change = e. An investor's value function is thus characterized by the following recursive equation:

$$V(X_t) = \max_{s_t^{dc}, s_t^l, \Xi_t} \left\{ u_t\left(c_t - k_\theta * 1\left\{\Xi_t \neq \Xi_{d,t}\right\} - k_s * 1\left\{s_t^{dc} \neq s_{d,t}\right\}\right) + (1 - m_{t+1})\beta E\left[V(X_{t+1}) \mid X_t\right] \right\},$$

subject to: (10), (11), (13), (14), (15), (17), (18), (20), and

$$s_t^{dc} * w_t + s_t^l = w_t - c_t + pen_{e,t} * s_t^{dc} * w_t * 1\{s_t^{dc} < 0\},$$
$$s_t^{dc} \leq limit_{e,t}.$$

Unemployment: $emp_t = U$**.** There are five sources of uncertainty when decisions are made at time t: mortality occurring with probability = m_{t+1}; ε_{t+1} = asset return shocks; possibility of becoming employed based on transition matrix; next-period labor income shocks conditional on becoming employed = η_{t+1}^U; the type of future employer after a job change = e.

$$V(X_t) = \max_{s_t^{dc}, s_t^l, \Xi_t} \left\{ u_t\left(c_t - k_\theta * 1\left\{\Xi_t \neq \Xi_{d,t}\right\}\right) + (1 - m_{t+1})\beta E\left[V(X_{t+1}) \mid X_t\right] \right\},$$

subject to: (12), (13), (14), (17), (18), (19), (20), and

$$s_t^l = ui_t - c_t - s_t^{dc} * w_t * (1 - pen_t),$$

$$s_t^{dc} \leq 0.$$

4.2 Estimation

We estimate the model parameters in two stages, as is common practice in life cycle model (e.g. Gourinchas and Parker 2002). The first-stage consists of setting parameters outside of the model based on auxiliary estimation and prior literature. The second stage consists of using Simulated Method of Moments to estimate the model's four preference parameters: the intertemporal discount factor (β), risk-aversion (σ), and the two adjustment costs (k_θ and k_s).

4.2.1 First-Stage Parameter Estimation

We set the risk-free rate, R_f, to be constant 2% and use a constant equity premium, $R_t^S - R_f$, of 4.5%, which is the average of the calibrations in Cocco et al. (2005) and Catherine (2020). Following Catherine (2020), we use a standard deviation of log returns, σ_s, of 19% and assume asset returns are uncorrelated with labor income. We use data from the Survey of Income Programs and Participation (SIPP) to estimate parameters of the labor income process and transition probabilities following Choukhmane (2021). The remaining first-stage parameters are estimated directly. Additional details on this first-stage estimation are provided in Appendix E.

4.2.2 Second-Stage Parameter Estimation

The four second-stage parameters are estimated using SMM, which minimizes the distance between model-simulated and data moments.

Empirical moments. We use 22 empirical moments in total. First, we use the average participation rates between tenures of 0 and 6 years for the control and treatment groups in our first quasi-experiments. This gives 14 moments, which are reported in Panel A of Figure 4. Second, we use the distribution of contribution rates among investors in our sample during their first year of tenure. We calculate the fraction of workers contributing the following fractions of their income: 0%, 3%, 6%, or 10% and above. We do this in two regimes: workers under the opt-in regime

and workers under an auto-enrollment regime with 3%.[24] As described below, this second set of moments of 8 moments is necessary to identify time preferences.

Model moments. To calculate the model counterparts to our first set of empirical moments, we perform our first quasi-experiment in our model. Specifically, we simulate 5,000 investors under two scenarios. First, investors are assigned to an employer with auto-enrollment into a money market fund at the start of every job. Secondly,at one randomly chosen job-transition, investors are randomly assigned (unexpectedly) to an employer with auto-enrollment into a fund with a 100% stock share. We then calculate the participation rates among these two groups of investors, which correspond to our treatment and control groups, within their first six years of tenure at the new jobs. To calculate the empirical counterparts to our second set of empirical moments, we again simulate 5,000 investors under two scenarios. First, investors face a 0% default contribution rate at the start of every job ("opt-in"). Second, investors face a 3% default contribution rate at one randomly chosen jobs. When then calculate the distribution of contribute rates across these investors.

Weighting matrix. We use the identify weighting matrix in our preliminary estimation following Altonji and Segal (1996).

4.2.3 Identification of Second-Stage Parameters

The four preference parameters in our model are jointly estimated. In what follows we provide some brief intuition for how parameters are identified.

Time preferences. The model has two parameters governing time preferences: the discount factor β and the inverse of relative risk aversion σ^{-1}. The fact that risk aversion controls time preferences and risk is a feature of time-separable CRRA preferences. Identification of these preference parameters comes from the distribution of contribution rates in the two regimes.

Risk preferences. In addition to affecting savings decisions, risk aversion also affects agents desire for stock market participation. Thus, risk aversion, σ, is identified from the levels of participation in treatment and control groups within our quasi-experiment. In the limit of extremely high risk-aversion, we would expect few investors to participate in the stock market in either groups.

Portfolio adjustment cost. The portfolio adjustment cost is identified by examining the number

[24] These moments are calculated on a sub-sample of our main data that is used by Choukhmane (2021). In particular, it consistent of 34 firms that: (i) have a 3% initial auto-enrollment default contribution rate with no auto-escalation feature; (ii) offer a 50% employer match contribution up to 6% of income; (iii) the exact date of auto-enrollment is available. Within these firms, we keep only investors hired within 12 months of the adoption of auto-enrollment.

of investors who deviate from the default asset allocation, in particular how this varies with tenure. In the limit of an infinite adjustment cost, we would expect all investors to stick with the default.

Contribution adjustment cost. The contribution adjustment cost is identified by examining the bunching of investors around the contribution rate into which they are defaulted. There are two sets of investors: those who have a default of 0% and those with a default of 3%. These investors have similar preferences, so the extent to which they bunch around the default determines the size of this adjustment cost. This identification strategy follows Choukhmane (2021).

4.3 Estimation Results

In a preliminary estimation, we find $\beta = 0.95$, $\sigma = 3.6$, $k_\theta = \$450$, and $k_s = \$450$. Our estimate of the discount factor is typical in the literature on life cycle consumption-savings model and portfolio choice (e.g. Gourinchas and Parker 2002). Our estimate of relative risk aversion is much lower than typical estimates in existing literature, as emphasized by Figure 1. This highlights the value of quasi-experimental variation using panel data to recover risk preferences in the presence of choice frictions. We find a similar value of the contribution adjustment cost to Choukhmane (2021). Our estimate of the portfolio adjustment cost is relatively low compared to the typical estimates required to get standard life cycle models to fit observed participation (e.g. Gomes 2020; Catherine 2020).

4.4 Model Fit

Figure 10 plots the fit of the model with respect to contribution rates. As evident from the figure, the model replicates this data quite well. Figure 11 plots the fit of our model with respect to the quasi-experiment in Figure 4. Our model can match the results of this quasi-experiment almost exactly, although we do not currently have standard errors.

5 Conclusion

This paper identifies investors' risk preferences in the presence of choice frictions. Although it is difficult to do so in general, separating preferences from frictions is important for positive reasons, such as distinguishing between competing economic models, but also for normative reasons, such

as assessing the impact of interventions designed to increase stock market participation on household welfare. For example, consider the Pension Protection Act of 2006, which led to the rapid growth of target date funds as the default asset in retirement savings plans. If non-participation in risky assets is primarily driven by frictions, then this trend is likely desirable for households. On the other hand, if non-participation mostly reflects preferences, then the welfare implications of this nudge towards the stock market are more ambiguous.

This paper has two main results. First, we estimate that absent frictions, 95% of investors would participate in the stock market with a stock share of 75%, which declines over the life cycle. Secondly, we show our evidence is consistent with a standard model of investor preferences, more specifically a constant coefficient of relative risk aversion of 3.6. Collectively, these results suggest investors' preferences in our sample are consistent with the predictions of textbook theories (e.g. Campbell and Viceira 2001), once we take into account choice frictions. As a result, our findings provide limited support for explanations of limited household stock market participation based on first-order risk aversion, such as loss aversion.

More broadly, we hope our analysis illustrates the caution needed when estimating behavioral models using cross-sectional moments. To the extent that frictions are present and impact choices, these frictions will obscure the mapping between observed choices and agents' underlying preferences. More constructively, our paper highlights how the use of quasi-experimental variation can help overcome this identification problem and provide better targets for testing different economic theories of investor behavior.

References

Abel, Andrew B., Janice Eberly, and Stavros Panageas (2013), "Optimal Inattention to the Stock Market With Information Costs and Transactions Costs." *Econometrica*, 81, 1455–1481.

Allcott, Hunt and Dmitry Taubinsky (2015), "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market." *American Economic Review*, 105, 2501–2538.

Altonji, Joseph G and Lewis M Segal (1996), "Small-sample bias in gmm estimation of covariance structures." *Journal of Business & Economic Statistics*, 14, 353–366.

Ameriks, John and Stephen P Zeldes (2004), "How Do Household Portfolio Shares Vary With Age?" *Working Paper*.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai (2020), "Sources of inaction in household finance: Evidence from the danish mortgage market." *American Economic Review*, 110, 3184–3230.

Andrews, Isaiah and Jesse M. Shapiro (2020), "A Model of Scientific Communication." *Working Paper*.

Barberis, Nicholas, Ming Huang, and Richard H. Thaler (2006), "Individual preferences, monetary gambles, and stock market participation: A case for narrow framing." *American Economic Review*, 96, 1069–1090.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein (2007), "Portfolio choice over the life-cycle when the stock and labor markets are cointegrated." *Journal of Finance*, 62, 2123–2167.

Bernheim, B. Douglas and Antonio Rangel (2009), "Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics." *Quarterly Journal of Economics*, 124, 51–104.

Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and William L. Skimmyhorn (2018), "Borrowing to Save? The Impact of Automatic Enrollment on Debt."

Blumenstock, Joshua, Michael Callen, and Tarek Ghani (2018), "Why do defaults affect behavior? Experimental evidence from Afghanistan." *American Economic Review*, 108, 2868–2901.

Briggs, Joseph, David Cesarini, Erik Lindqvist, and Robert Östling (2021), "Windfall gains and stock market participation." *Journal of Financial Economics*, 139, 57–83.

Brunnermeier, Markus K. and Stefan Nagel (2008), "Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals' asset allocation." *American Economic Review*, 98, 713–736.

Calvet, Laurent E., John Y. Campbell, Francisco J. Gomes, and Paolo Sodini (2019), "The Cross-Section of Household Preferences." *Working Paper*.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini (2009), "Fight or flight? Portfolio rebalancing by individual investors." *The Quarterly Journal of Economics*, 124, 301–348.

Campbell, John Y. (2006), "Household finance." *Journal of Finance*, 61, 1553–1604.

Campbell, John Y. (2016), "Restoring rational choice: The challenge of consumer financial regulation." *American Economic Review: Papers & Proceedings*, 106, 1–30.

Campbell, John Y. and Luis M. Viceira (2001), *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*.

Carroll, Christopher D (2020), *Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems*, volume 7602.

Catherine, Sylvain (2020), "Countercyclical Income Risk and Portfolio Choices over the Life-Cycle." *Working Paper*.

Chapman, David A. and Valery Polkovnichenko (2009), "First-Order risk aversion, heterogeneity, and asset market outcomes." *Journal of Finance*, 64, 1863–1887.

Chetty, Raj (2006), "A new method of estimating risk aversion." *American Economic Review*, 96, 1821–1834.

Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen (2014), "Active vs. Passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark." *Quarterly Journal of Economics*, 129, 1141–01219.

Choukhmane, Taha (2021), "Default Options and Retirement Saving Dynamics." *Working Paper*, 1–78.

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

Dahlquist, Magnus, Ofer Setty, and Roine Vestman (2018), "On the Asset Allocation of a Default Pension Fund." *Journal of Finance*, 73, 1893–1936.

Deaton, Angus S and Christina H Paxson (1994), "Saving, Growth, and Aging in Taiwan." *Working Paper*.

Duarte, Victor, Julia Fonseca, Jonathan A. Parker, and Aaron Goodman (2021), "Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle." *Working Paper*.

Enke, Benjamin and Thomas Graeber (2020), "Cognitive Uncertainty." *Working Paper*.

Epstein, Larry and Tan Wang (1994), "Intertemporal Asset Pricing under Knightian Uncertainty." *Econometrica*, 62, 283–322.

Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso (2017), "Asset Market Participation and Portfolio Choice over the Life-Cycle." *Journal of Finance*, 72, 705–750.

Gabaix, Xavier (2019), "A Behavioral New Keynesian Model."

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen P. Utkus (2021), "Five Facts About Beliefs and Portfolios." *American Economic Review*, 115, 1481–1522.

Goldin, Jacob and Daniel Reck (2020), "Revealed-preference analysis with framing effects." *Journal of Political Economy*, 128, 2759–2795.

Gomes, Francisco and Alexander Michaelides (2005), "Optimal life-cycle asset allocation: Understanding the empirical evidence." *Journal of Finance*, 60, 869–904.

Gomes, Francisco J. (2005), "Portfolio choice and trading volume with loss-averse investors." *Journal of Business*, 78, 675–706.

Gomes, Francisco J. (2020), "Portfolio Choice over the Life Cycle: A Survey." *Annual Review of Financial Economics*, 12, 277–304.

Gomes, Francisco J., Michael Haliassos, and Tarun Ramadorai (2020), "Household finance." *Working Paper*.

Gomes, Francisco J., Alexander Michaelides, and Yuxin Zhang (2021), "Tactical Target Date Funds." *Management Science*.

Gourinchas, Pierre Olivier and Jonathan A. Parker (2002), "Consumption over the life cycle." *Econometrica*, 70, 47–89.

Guiso, Luigi, Michael Haliassos, and Tullio Jappelli (2002), *Household Portfolios*. MIT Press, Cambridge, MA.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2008), "Trusting the stock market." *Journal of Finance*, 63, 2557–2600.

Guvenen, Fatih (2009), "An empirical investigation of labor income processes." *Review of Economic dynamics*, 12, 58–79.

Haliassos, Michael and Alexander Michaelides (2003), "Portfolio Choice and Liquidity Constraints." *International Economic Review*, 44, 143–177.

Hansen, Lars Peter and Kenneth J Singleton (1982), "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica*, 50, 1269–1286.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Jae Song (2020), "Technological Innovation and Labor Income Risk." *Working Paper*.

Laibson, David, Peter Maxted, and Benjamin Moll (2020), "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy." *Working Paper*, 1–57.

Low, Hamish, Costas Meghir, and Luigi Pistaferri (2010), "Wage risk and employment risk over the life cycle." *American Economic Review*, 100, 1432–67.

Lusardi, Annamaria, Pierre Carl Michaud, and Olivia S. Mitchell (2017), "Optimal financial knowledge and wealth inequality." *Journal of Political Economy*, 125, 431–477.

Madrian, Brigitte C. and Dennis F. Shea (2001), "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *The Quarterly Journal of Economics*, 4, 1149–1186.

Mankiw, N. Gregory and Stephen P Zeldes (1991), "The consumption of stockholders and nonstockholders." *Journal of Financial Economics*, 29, 97–112.

Masatlioglu, Yusufcan and Efe A Ok (2005), "Rational choice with status quo bias." *Journal of Economic Theory*, 121, 1–29.

Massa, Massimo, Rabih Moussawi, and Andrei Simonov (2021), "The Unintended Consequences of Investing for the Long Run: Evidence from Target Date Funds." *Working Paper*.

Meeuwis, Maarten (2020), "Wealth Fluctuations and Risk Preferences: Evidence from U.S. Investor Portfolios." *Working Paper*.

Mehra, Rajnish and Edward C. Prescott (1985), "The Equity Premium A Puzzle." *Journal of Monetary Economics*, 15, 145–161.

Merton, Robert C. (1969), "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case." *Review of Economics and Statistics*, 51, 247–257.

Pagel, Michaela (2018), "A News-Utility Theory for Inattention and Delegation in Portfolio Choice." *Econometrica*, 86, 491–522.

Parker, Jonathan A., Antoinette Schoar, and Yang Sun (2020), "Retail Financial Innovation and Stock Market Dynamics : The Case of Target Date Funds." *Working Paper*.

Rabin, Matthew (2000), "Risk Aversion and Expected Utility Theory: A Calibration Theorem." *Econometrica*, 68, 1281–1292.

Reher, Michael and Stanislav Sokolinski (2021), "Automation and Inequality in Wealth Management." *Working Paper*.

Rubin, Donald B. (1978), "Bayesian Inference for Causal Effects: The Role of Randomization." *Annals of Statistics*, 6, 34–58.

Straub, Ludwig (2019), "Consumption, Savings, and the Distribution of Permanent Income." *Working Paper*, 1–85.

Vissing-Jørgensen, Annette (2002), "Limited asset market participation and the elasticity of intertemporal substitution." *Journal of Political Economy*, 110, 825–853.

FIGURES AND TABLES

Figure 1. Estimates of Relative Risk Aversion in Life Cycle Portfolio Choice

This figure plots the coefficient of relative risk aversion that is calibrated or estimated in papers that solve dynamic life cycle portfolio choice models with CRRA or Epstein-Zin preferences. We also include the type of the friction associated with stock market participation in the model, distinguishing between one-time and per-period costs. If the paper has multiple estimations, we include the paper's main estimates. An exception is Calvet, Campbell, Gomes, and Sodini (2019) who estimate risk-aversion across investors; we include their average estimate. Merton Share refers to the portfolio share implied by the Merton (1969) model with a risk-premium of 5.5% and standard deviation of 16%.

Paper	Friction	Estimate	Merton Share
Fagereng et al. (2017)	per-period cost	14.4	15%
Dahlquist et al. (2018)	one-time cost	14	15%
Fagereng et al. (2017)	per-period cost	11.8	18%
Catherine (2020)	none	11.6	19%
Fagereng et al. (2017)	per-period cost	11	20%
Cocco et al. (2005)	none	10	21%
Reher and Sokolinski (2021)	per-period cost	9.1	24%
Catherine (2020)	per-period cost	8.2	26%
Catherine (2020)	none	6.9	31%
Catherine (2020)	per-period cost	6.2	35%
Calvet et al. (2020)	none	5.24	41%
Gomes and Michaelides (2005)	one-time cost	5	43%
Campanale et al. (2015)	one-time cost	5	43%
Benzoni et al. (2007)	none	5	43%
Gomes et al. (2009)	one-time cost	4	54%
Briggs et al. (2021)	one-time cost	3.1	69%
Briggs et al. (2021)	one-time cost	2.4	90%
Average		**7.82**	**27%**

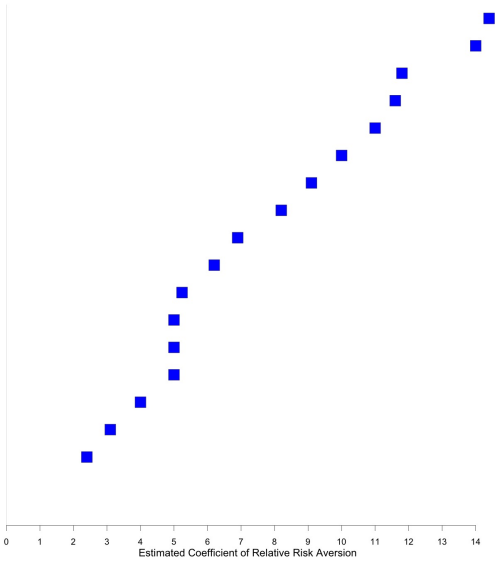


Figure 2. Preferences vs. Frictions

This figure plots the life cycle of participation for different parameterizations of the model in Appendix A, using a discount factor of $\beta = 0.96$ in all simulations. We also plot the life-cycle of participation from the SCF 1989-2019, identified using the methodology in Deaton and Paxson (1994), with 95% confidence intervals. Each model simulation consists of 1,000 investors.

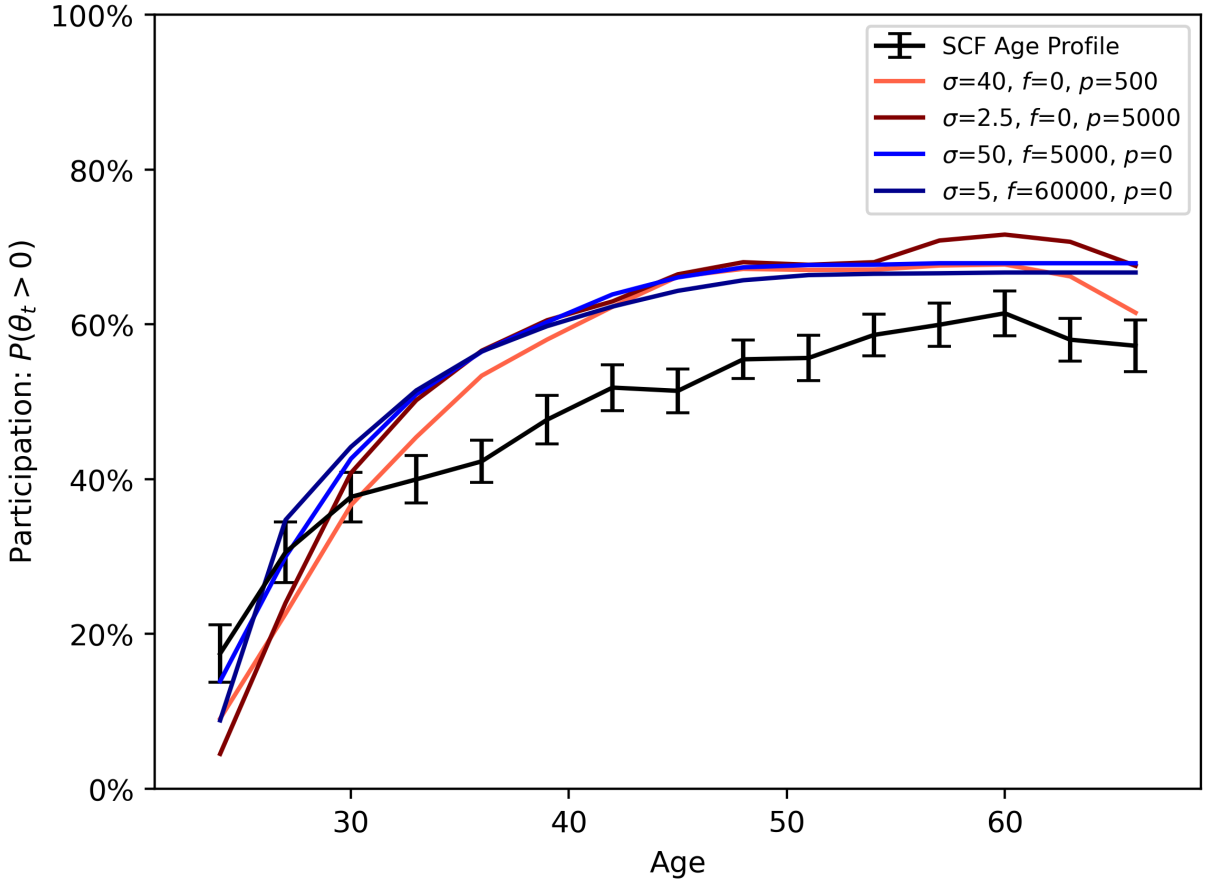


Table 1. Summary Statistics

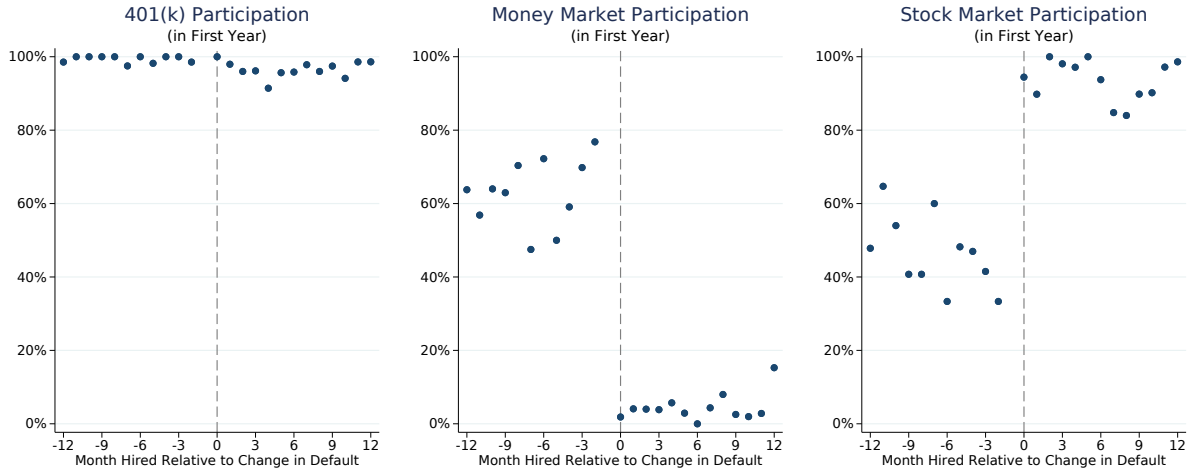
This table provides summary statistics on our sample versus the SCF. We calculate summary statistics on our sample using the entire universe of investors provided by out data provider, where each observation is at the investor/employee-year level. We calculate summary statistics from the 2007, 2010, 2013, and 2016 SCF waves, where we adjust survey weights such that they assign equal weights to each survey wave. SCF equity shares are calculated by dividing total equity holdings by total financial wealth and participation based on whether this is positive. Wage income, financial wealth, and retirement wealth from the SCF are divided by the number of adults in the household. Retirement wealth is in the SCF is defined as the sum of total quasi-liquid retirement accounts, including IRAs, thrift accounts, future pensions, and currently received benefits. We define SCF investors as being eligible for a retirement account if they report having access to a retirement account and/or they report assets in one.

	SCF 2007-2016				Our Sample 2006-2017	
	All Households $N = 16,187$		Retirement Account Eligible $N = 7,639$		$N = 18,398,750$	
	Mean	Median	Mean	Median	Mean	Median
Age	49.62	49.00	45.24	46.00	41.59	41.00
Wage Income Per Adult	77,193.50	44,577.58	100,135.55	67,980.80		33,854.40
Financial Wealth	249,413.12	17,521.30	300,179.88	45,368.47		
Retirement Wealth	97,580.51	1,178.25	141,844.65	24,688.21	49,420.93	5,554.77
Ratio of Retirement to Financial Wealth	0.33	0.08	0.53	0.61		
Equity Share (Financial Wealth)	0.23	0.03	0.34	0.29		
Equity Share (Retirement Wealth)	0.27	0.00	0.42	0.40	0.53	0.73
Ratio of Equity Holdings in Retirement to Total	0.38	0.00	0.62	0.95		
Participation (Financial Wealth)	0.52	1.00	0.76	1.00		
Participation (Retirement Wealth)	0.46	0.00	0.72	1.00	0.68	1.00
Participation (Outside Retirement)	0.15	0.00	0.16	0.00		
Participation (Only Outside Retirement)	0.05	0.00	0.03	0.00		

Figure 3. Identifying Variation

This figure plots the variation in our two quasi-experiments. In Panel A, we compare the portfolio choices and 401(k) participation of employees hired within 12 months before and 12 months after 6 firms change the default asset allocation in their auto-enrollment 401(k) plans. The control group is 1,086 employees hired before the change, who are defaulted into a money market fund (i.e. no stock market exposure), and the treatment group are 1,321 investors hired *at the same firms* after the change, who are defaulted into a TDF (i.e. has stock market exposure). In Panel B, we compare the portfolio choices and 401(k) participation of employees hired within 12 months before and after 191 firms change their 401(k) plans from an opt-in regime to automatic-enrollment with a TDF as the default asset allocation. The control group is 40,337 employees hired before the change under the opt-in regime, while the treatment group is 52,400 employees hired after the change and automatically-enrolled into a TDF. See Section 2.3 for more details.

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

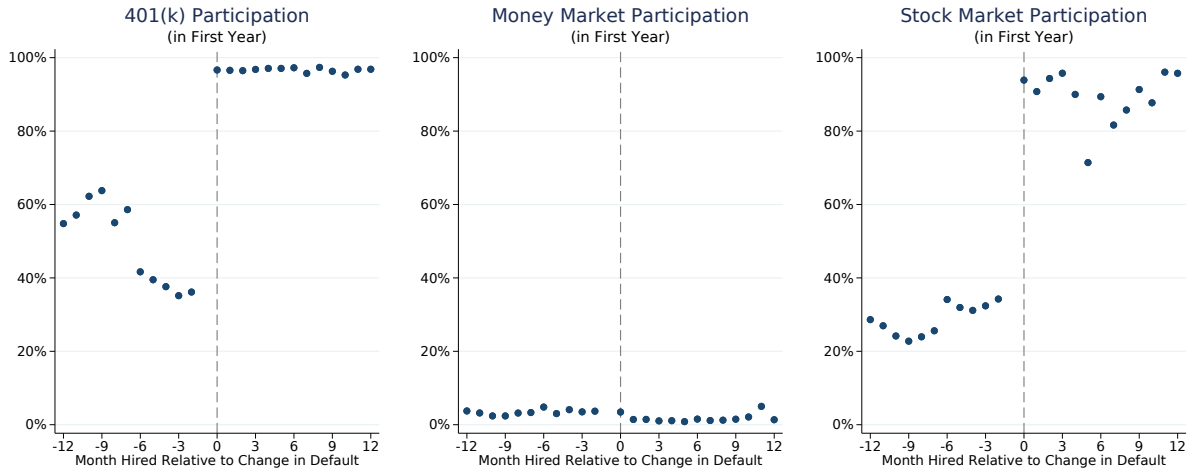
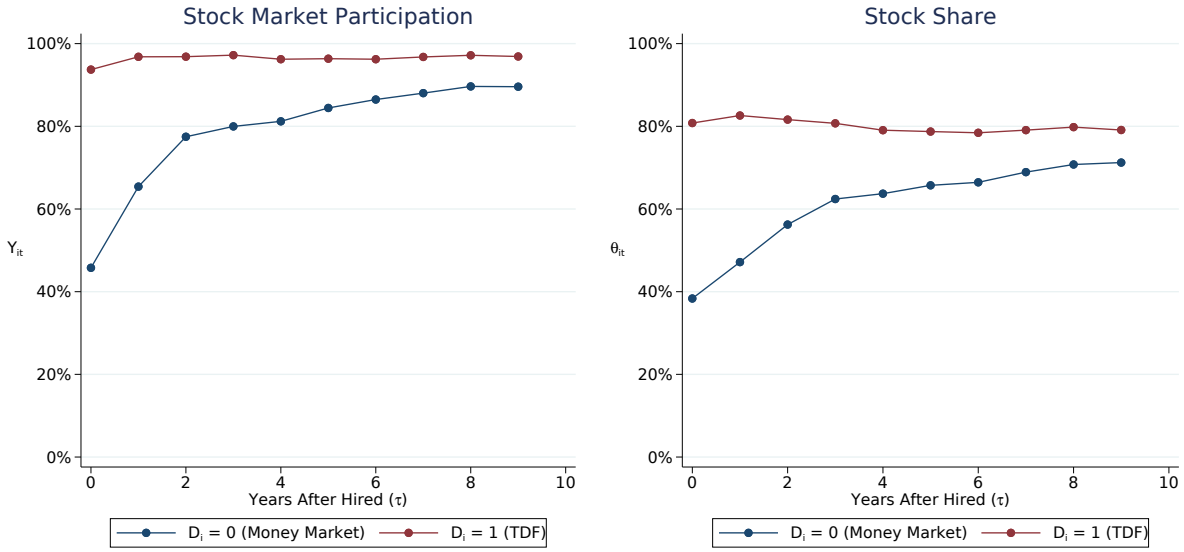


Figure 4. Observed Portfolio Choice Response

This figure plots the observed portfolio responses for investors who are defaulted into two groups. In Panel A, we compare investors defaulted into an automatic-enrollment 401(k) plan with two different default asset allocations: a money market fund and a TDF. Panel B does the same comparison between investors defaulted into an opt-in 401(k) plan and those defaulted into an automatic enrollment 401(k) plan with a TDF as the default asset. See Figure A4 for the average choices in each sample by default.

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

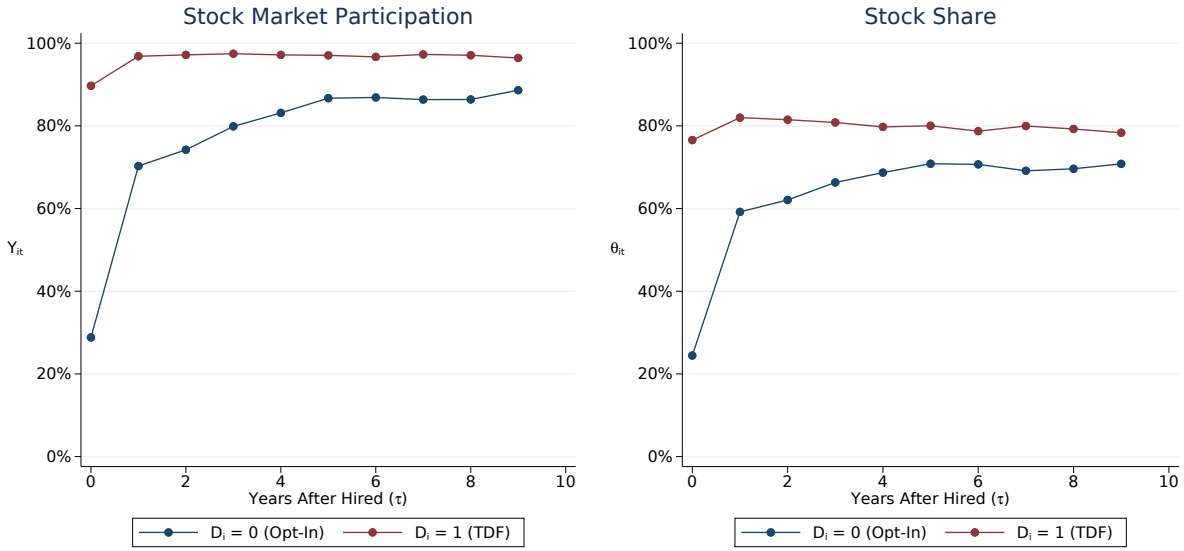


Table 2. Observed Portfolio Choice Response: Regression

This table displays the regression results that complement Figure 4, in which we regress investors' observed choices onto an indicator for whether they are in the treatment group and thus have a default asset allocation with stock market exposure. Panel A displays the results comparing investors defaulted into an automatic-enrollment 401(k) plan with a money market fund default and those with a and a TDF default. Panel B displays analogous results comparing investors defaulted into an opt-in 401(k) plan with those defaulted into an automatic enrollment 401(k) plan with a TDF as the default asset. In both panels, two-way clustered standard errors by firm and year are shown in parenthesis.

Panel A: Money Market to TDF Sample

	Participation: Y_{it}				Stock Share: θ_{it}			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	56.35	56.35	56.95	56.95	37.21	37.21	38.28	38.28
	(10.57)	(10.57)	(8.600)	(8.600)	(9.392)	(9.392)	(8.004)	(8.004)
Default Has Stocks: D_i	20.06	20.06	19.68	19.68	21.62	21.62	20.93	20.93
	(5.596)	(5.596)	(5.795)	(5.795)	(5.239)	(5.239)	(5.215)	(5.215)
Tenure Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Firm Fixed Effects			✓	✓			✓	✓
Firm and Year Clustering	✓	✓	✓	✓	✓	✓	✓	✓
Total Observations	12650	12650	12650	12650	12650	12650	12650	12650
Adjusted R-Squared	0.124	0.124	0.149	0.149	0.130	0.130	0.146	0.146

Panel B: Opt-In to TDF Sample

	Participation: Y_{it}				Stock Share: θ_{it}			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	65.28	65.28	66.97	66.97	54.30	54.30	55.46	55.46
	(2.913)	(3.113)	(1.688)	(1.688)	(2.474)	(2.648)	(1.337)	(1.337)
Default Has Stocks: D_i	30.13	30.13	27.12	27.12	25.70	25.70	23.65	23.65
	(2.707)	(2.822)	(2.798)	(2.798)	(2.287)	(2.344)	(2.256)	(2.256)
Tenure Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Firm Fixed Effects			✓	✓			✓	✓
Firm and Year Clustering	✓	✓	✓	✓	✓	✓	✓	✓
Total Observations	263061	263061	263061	263061	263061	263061	263061	263061
Adjusted R-Squared	0.230	0.230	0.323	0.323	0.197	0.197	0.290	0.290

Figure 5. Bounding Population Preferences: Money Market to TDF Sample

This figure plots the same data as Figure 4, but indicates the non-parametric bounds on average preferences in the population based on Assumptions 1-4. See Section 3.2 for additional details.

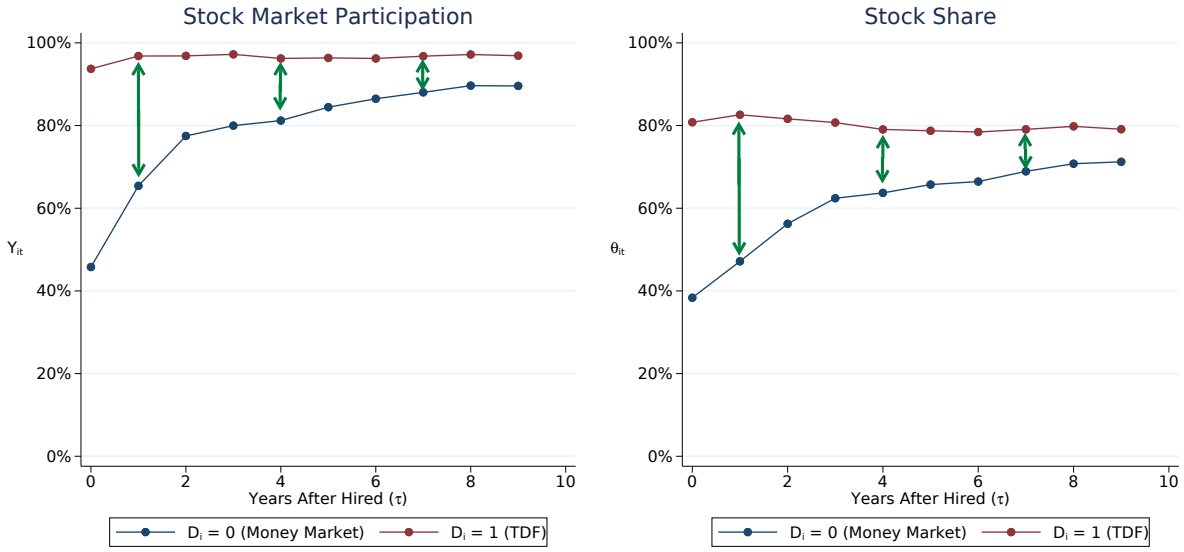
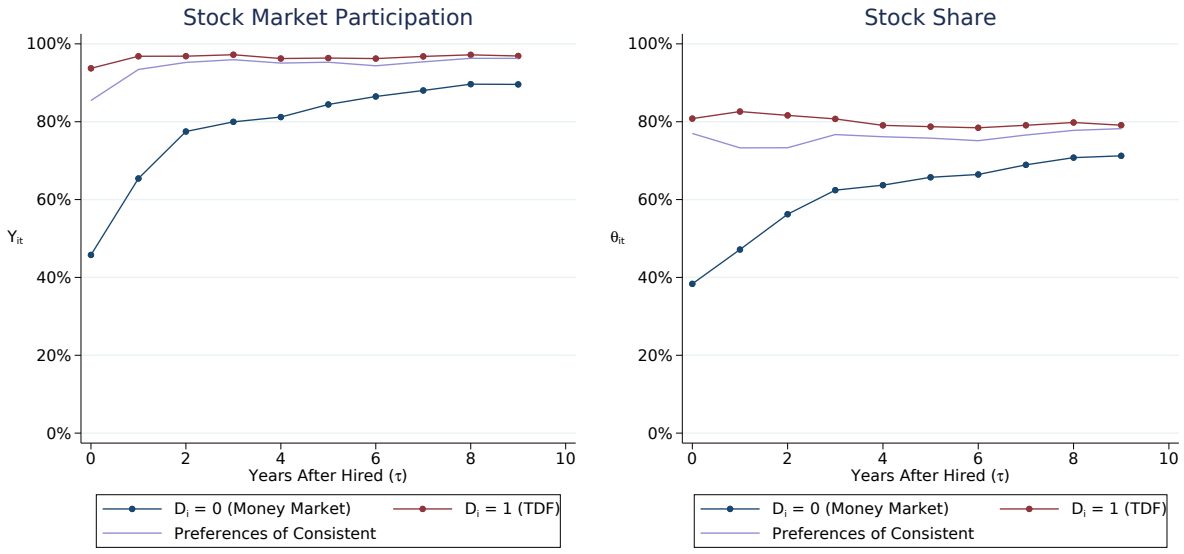


Figure 6. Preferences of Consistent Investors

This figure plots the same data as Figure 4, but includes our point estimates for the consistent investors under 1-4. See Section 3.2 for additional details. Under an additional assumption of 4′, this provides an estimate of the preferences of the entire population, as discussed in Section 3.4.

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

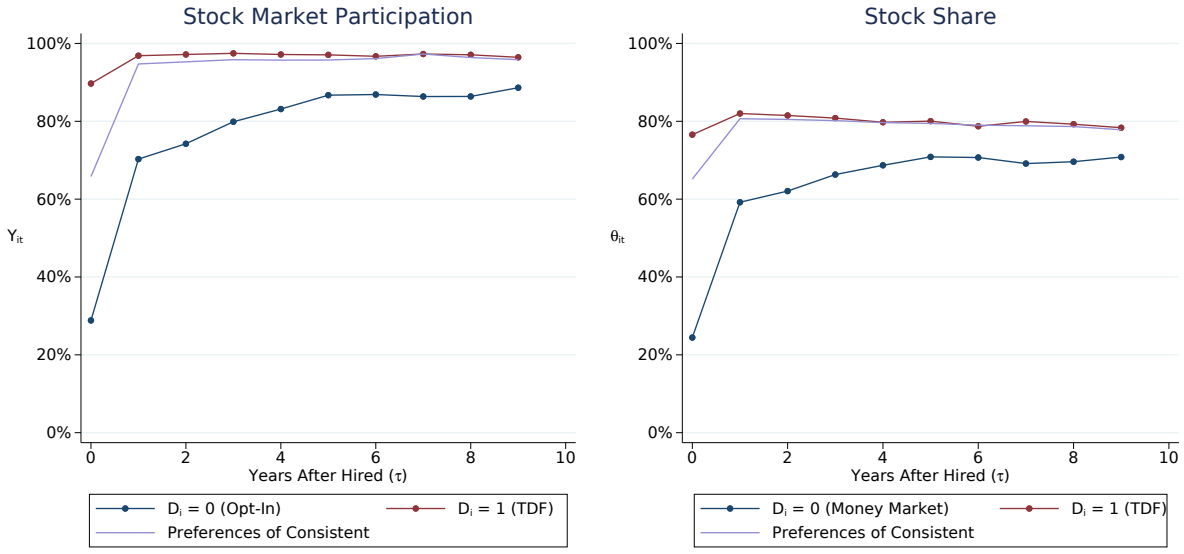
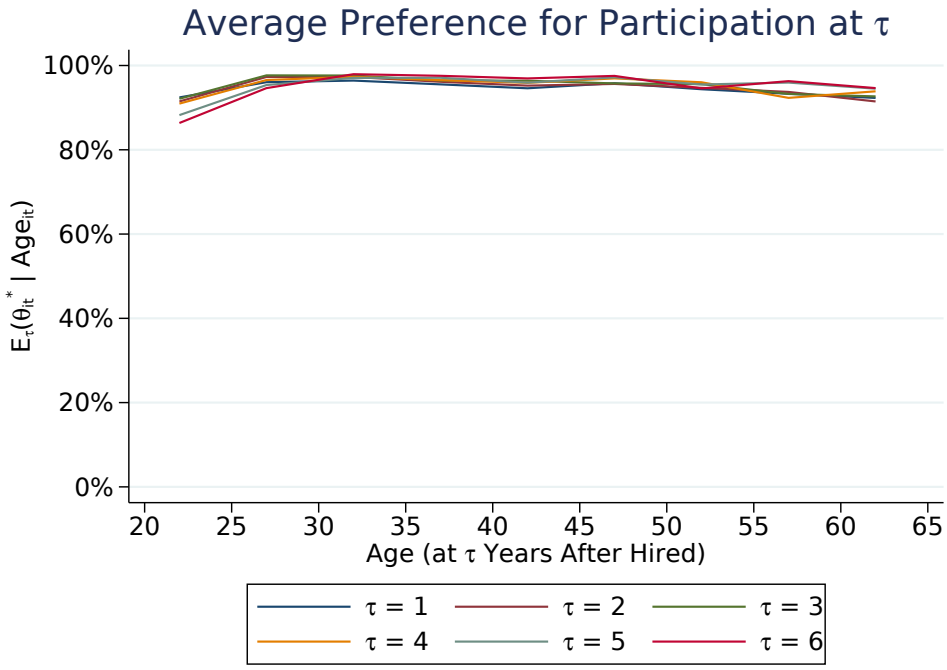


Figure 7. Participation Preferences vs. Choices over the Life Cycle: Opt-In to TDF Sample

This figure plots our estimate of investors' preferences for participation in Panel A in comparison to their observed choices in Panel B. For reference, Panel B also includes data from the SCF 2007 wave. In Panel A, each line corresponds to our estimates of preferences of investors at different tenures. See Section 3.5 for more details.

Panel A: Estimated Preferences



Panel B: Observed Choices

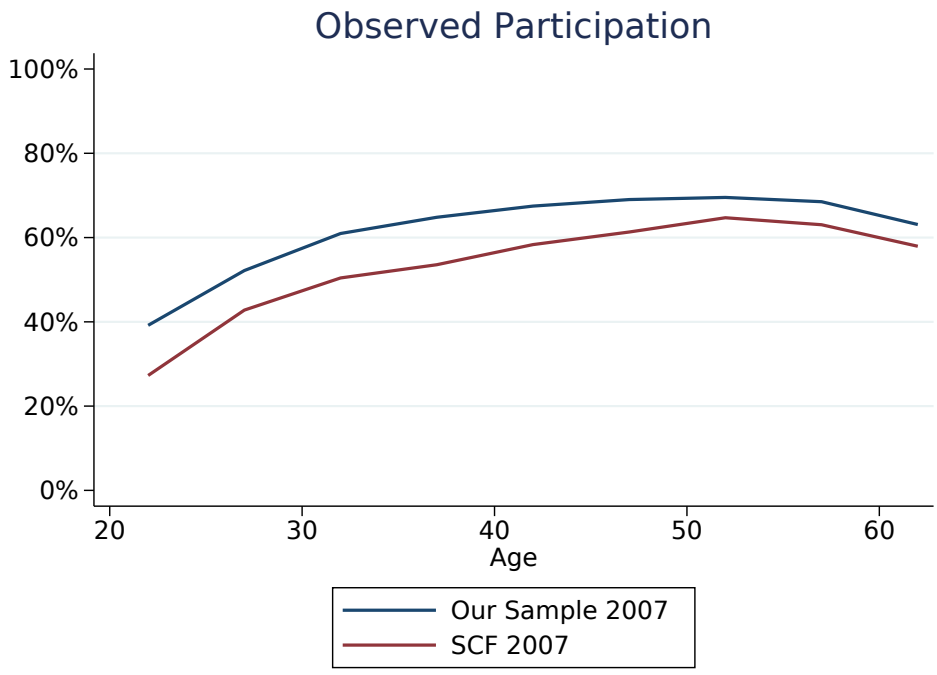
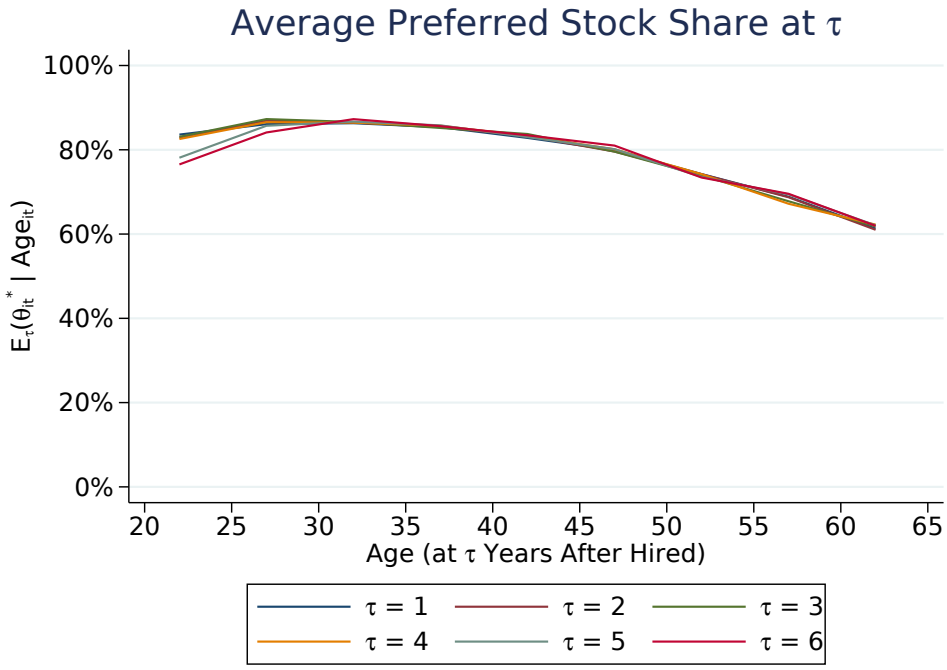


Figure 8. Stock Share Preferences vs. Choices over the Life Cycle: Opt-In to TDF Sample

This figure plots our estimate of investors' preferences for stock shares in Panel A in comparison to their observed choices in Panel B. For reference, Panel B also includes data from the SCF 2007 wave. In Panel A, each line corresponds to our estimates of preferences of investors at different tenures. See Section 3.5 for more details.

Panel A: Estimated Preferences



Panel B: Observed Choices

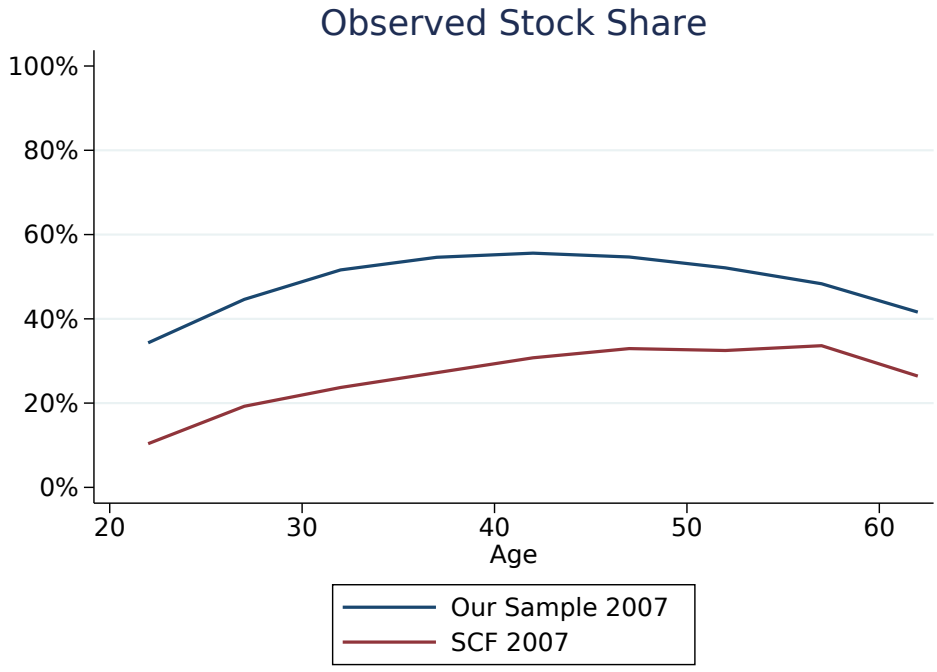
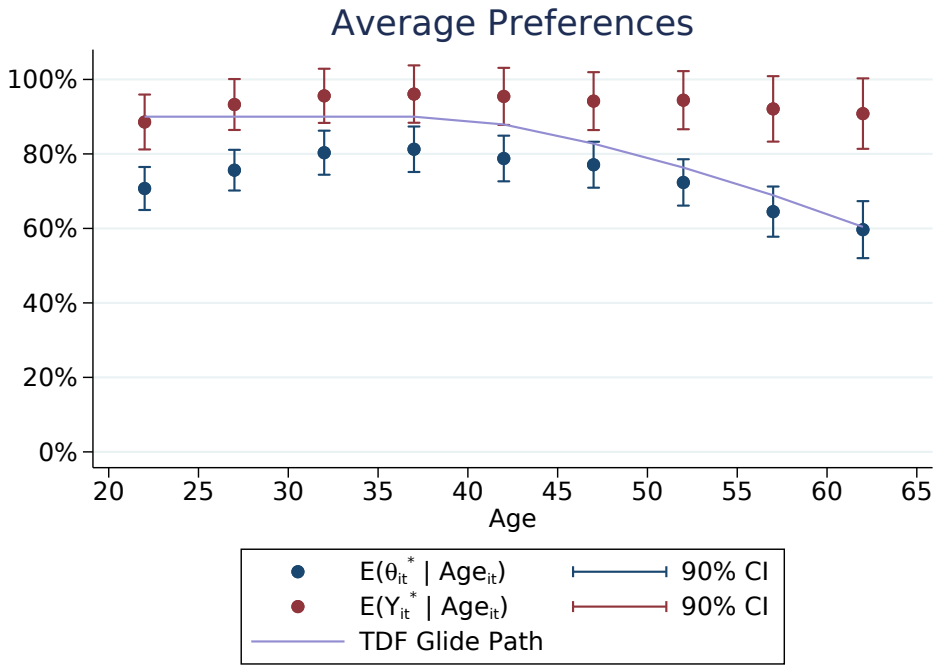


Figure 9. Average Estimated Preferences over the Life Cycle

This figure plots our estimate of investors' preferences for participation and stock shares over the life cycle. These preferences are calculating by integrating over the tenure-specific results in Figures 7 and 8. The plot also includes 90% confidence intervals and the glide-path for the TDF into which the treatment group is defaulted. Standard errors are clustered by employee in Panel A and firm in Panel B, since there are only 6 firms in Panel A.

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

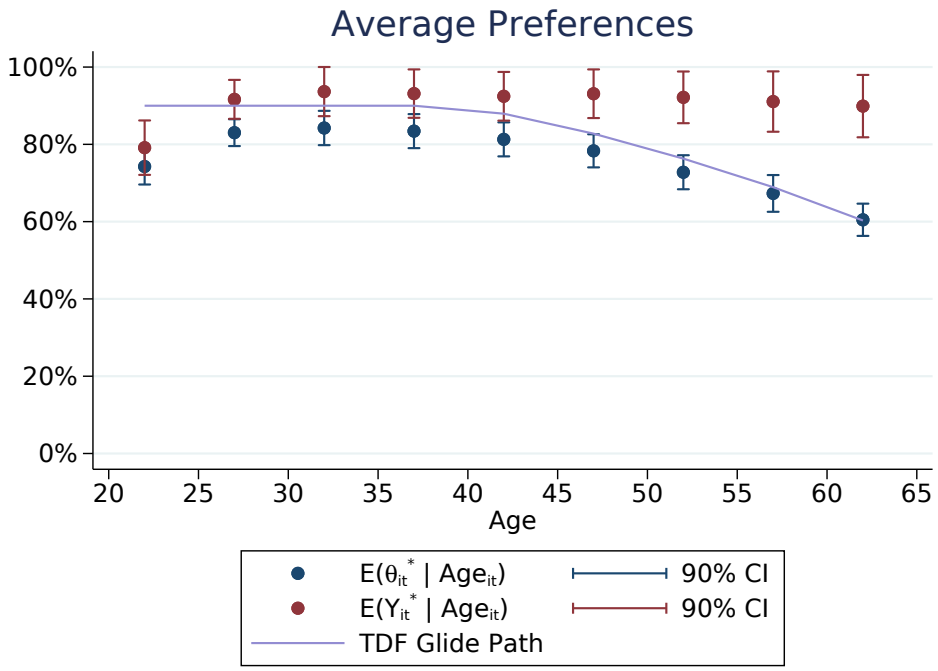


Figure 10. Model Fit: Contribution Rates

This figure presents a comparison of one set of targeted moments in our structural estimation between our data and model. The top (bottom) figure show contribution rates of investors hired 12 months before (after) the introduction of auto-enrollment for new hires. The model results are from a simulation of this within the model at our SMM estimates of preference parameters discussed in Section 4.2. The data results are discussed in Section 4.2.

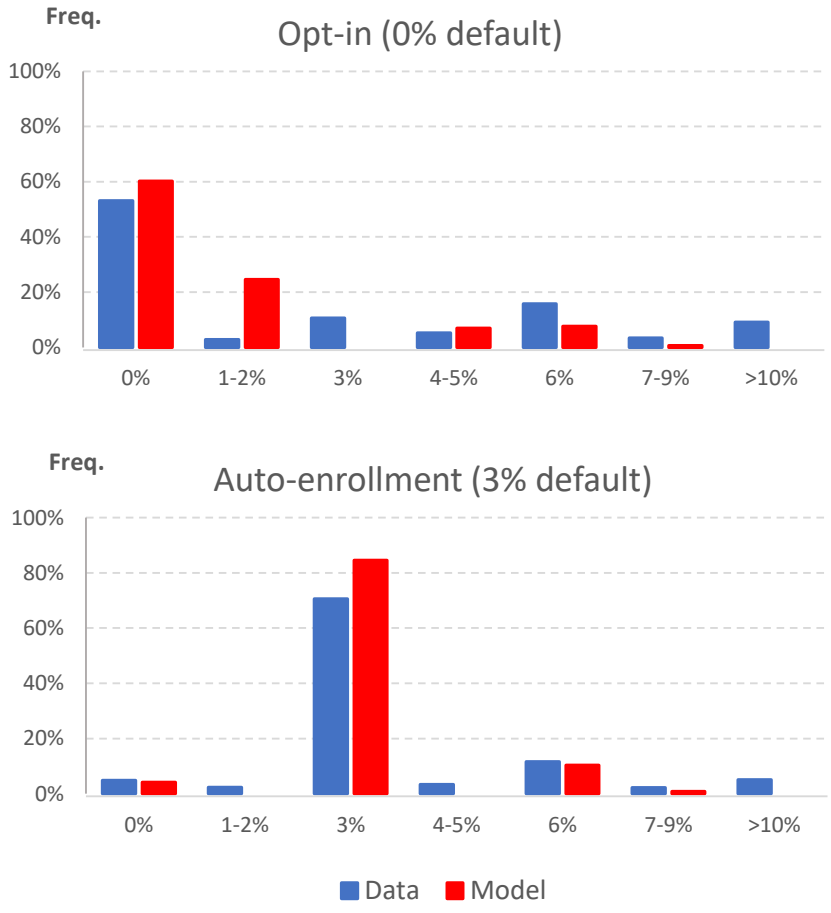
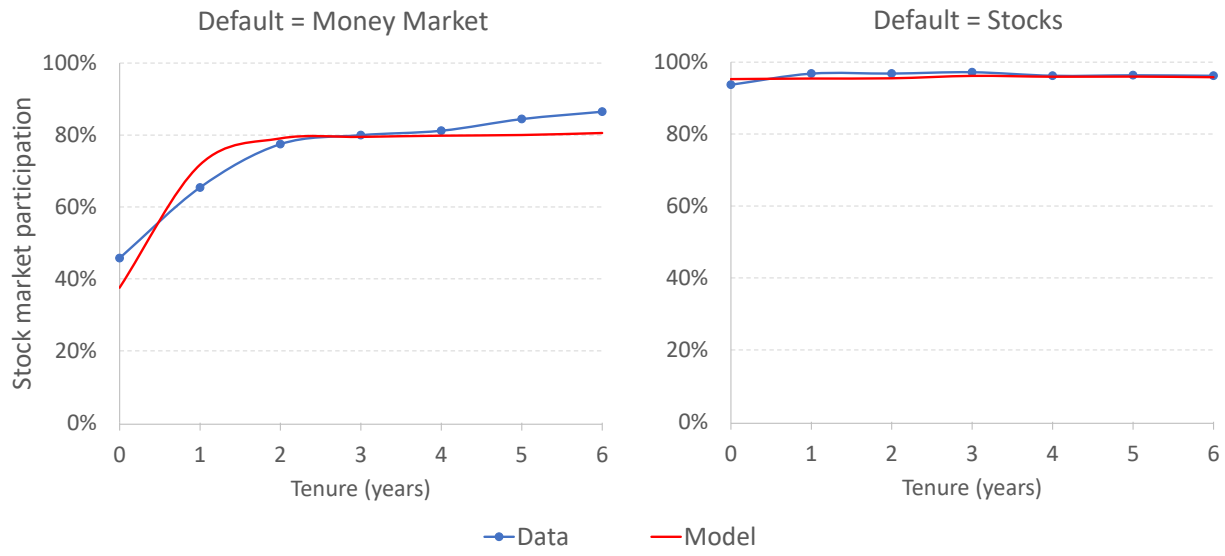


Figure 11. Model Fit: Quasi-Experimental Evidence

This figure presents a comparison of one set of targeted moments in our structural estimation between our data and model. The data moments in this figure correspond to the moments from our first quasi-experiment in Figure 4 Panel A for the first six years of tenure. The model results are from a simulation of this experiment within the model at our SMM estimates of preference parameters discussed in Section 4.2.



APPENDIX

Appendix A. Simple Life Cycle Model

A.1 Demographics

Investors are born at $t = 0$ and work T^w periods with their first retirement year at $t = T^w$. Each period corresponds to one year. Investors die with certainty at $t = \overline{T}$, at which point all of their resources are bequested, such that investor's last period in which they can consume is $t = \overline{T} - 1$. Before their certain death, investors survival probability is time-varying and denoted m_t, which is taken from the SSA. Denote investor's age as $a_t = t + a_0$, where a_0 is the age at which investors are born.

A.2 Preferences

Investors have time-separable expected utility preferences with a CRRA Bernoulli utility function over consumption. Denote investors' annualized time discount factor as β and their coefficient of relative risk aversion (or equivalently inverse of elasticity of intertemporal substitution) as σ. Per-period flow utility is adjusted for an equivalence scale such that it is equal to

$$u_t(c) = n_t * \frac{\left(c/n_t\right)^{1-\sigma}}{1-\sigma},$$

where n_t is taken from Lusardi et al. (2017).

A.3 Labor Income

While working, investors supply labor inelastically and earn an income w_t that is stochastic and exogenous. This income consists of a deterministic component that is cubic in age and a stochastic component that follows an AR(1) process with normal innovations:

$$\ln w_t = \delta_0 + \delta_1 a_t + \delta_2 a_t^2 + \delta_3 a_t^3 + \eta_t, \quad \eta_t = \rho \eta_{t-1} + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2). \tag{21}$$

Investors do not receive any labor income at the time of death. Additionally, when investors are born at $t = 0$, the distribution of η is different to account for heterogeneity in the initial period income shock.

When $t \in [T^w, \overline{T} - 1]$, investors are retired and earn retirement benefits denoted by ss_t, which are set equal to 40% of the average annual wage among all working investors. Note that investors do not earn any retirement benefits at the time of death.

A.4 Financial Assets

There are two assets: (i) a risk-free bond that has a constant gross return of R_f per year; (ii) a risky asset that pays a stochastic IID gross return of R_t per year, where

$$\ln R_t = \ln R_f + \mu_s + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_s^2). \tag{22}$$

A.5 Savings Account

Investors have access to a liquid savings account in which they can invest any remaining labor income after consuming into any combination of the two available assets. Denote the dollar amount in this account as A_t. Investors choose the share of their wealth that is invested in the risky asset, denoted by θ_t, resulting in $1 - \theta_t$ allocated to the risk-free bond.

A.6 Frictions

When making portfolio choice decisions, investors face two costs. First, there is a per-period participation cost p, which is incurred when $\theta_t > 0$ as a utility cost. This cost is designed to capture the costs associated with maintaining an account to invest in the risky asset, in addition to any hassle costs.

Secondly, investors must incur a utility cost f to change their portfolio, which is designed to capture default effects in investor's behavior. Specifically, if investors choose $\theta_t \neq \theta_{d,t}$, where

$$\theta_{d,t} = \begin{cases} 0 & \text{if } t = 0, \\ \theta_{t-1} * \frac{R_t}{(1-\theta_{t-1})R_f + \theta_{t-1}R_t} & \text{else.} \end{cases} \tag{23}$$

they are required to pay a cost f in consumption units that is designed to capture the physical and opportunity costs associated with altering a portfolio allocation. We assume the default asset allocation in the first period is entirely risk-free bonds to capture the fact that investors are generally not born with exposure to stocks. Default asset allocations in later periods are equal to the asset allocation from the prior period, adjusting for return realizations. Additionally, we assume neither of these costs are incurred at time t if an investor dies at time t.

A.7 Optimization Problem

Investors face a dynamic optimization problem with four state variables: a_t = age; w_t = wage; W_t = liquid wealth; θ_{t-1} = prior-period portfolio share. Denote the vector of these variables as X_t. There are two controls: c_t = consumption and θ_t = stock share. In choosing these controls, investors are restricted from borrowing and engaging in any margin

trading (i.e. no short-selling or leveraged positions):

$$W_t \geq 0, \quad \theta_t \in [0,1]. \tag{24}$$

We now characterize the value function of an investor, $V(\cdot) : \mathbb{R}^4_+ \to \mathbb{R}$, in the two periods of their life.

Retirement period. In retirement, there are two sources of uncertainty when decisions are made at time t: ε_{t+1} = asset return shocks; mortality occurring with probability m_{t+1}. An investor's value function during retirement is thus characterized by the following recursive equation:[25]

$$V(X_t) = \max_{c_t>0,\theta_t} \left\{ u_t\left(c_t - f * 1\{\theta_t \neq \theta_{d,t}\} - p * 1\{\theta_t > 0\}\right) + (1-m_{t+1})\beta E\left[V(X_{t+1}) \mid X_t\right] \right\},$$

subject to: (22), (23), (24), and

$$A_t = W_t + ss_t - c_t,$$
$$W_t = A_{t-1}\left[(1-\theta_{t-1})R_f + \theta_{t-1}R_t\right],$$
$$V(a_{\overline{T}}, \cdot, \cdot, \cdot) = 0.$$

Working life. While working, there are three sources of uncertainty when decisions are made at time t: η_{t+1} = labor income shocks; ε_{t+1} = asset return shocks; mortality occurring with probability m_{t+1}. An investor's value function during working-life is thus characterized by the following recursive equation:

$$V(X_t) = \max_{c_t>0,\theta_t} \left\{ u_t\left(c_t - f * 1\{\theta_t \neq \theta_{d,t}\} - p * 1\{\theta_t > 0\}\right) + (1-m_{t+1})\beta E\left[V(X_{t+1}) \mid X_t\right] \right\},$$

subject to: (21), (22), (23), (24), and

$$A_t = W_t + w_t - c_t,$$
$$W_t = A_{t-1}\left[(1-\theta_{t-1})R_f + \theta_{t-1}R_t\right],$$
$$A_0 = W_0 = 0.$$

A.8 Parameterization

For the income process parameters (ρ and σ_v) we use the parameters estimated in Appendix E for our full model for agents that are in $emp_t = E$ (i.e. continuously employed). For demographic parameters, we set $a_0 = 21$, $T^w = 44$ so agents retire deterministically at 65, and $\overline{T} = 79$ so agents die at 100. We take m_t from the SSA as in Appendix E. We use the same parameters for asset returns as in Section 4. We use a $\beta = 0.99$ in Figure 2 and Figure A1.

[25] Note that in principle we could define a separate value function for retirement since the investor's policies are independent of w_t, but do not for notational simplicity.

Appendix B. Proofs and Derivations

Proof of Proposition 1. By the law of iterated expectations, we obtain

$$E_\tau(Y_{it}^*) = E_\tau(Y_{it}^* \mid C_{it}=1)P_\tau(C_{it}=1) + E_\tau(Y_{it}^* \mid C_{it}=0)P_\tau(C_{it}=0).$$

Using the fact that Y_{it}^* is bounded between zero and one, the previous equation implies

$$E_\tau(Y_{it}^*) \in \left[E_\tau(Y_{it}^* \mid C_{it}=1)P_\tau(C_{it}=1), E_\tau(Y_{it}^* \mid C_{it}=1)P_\tau(C_{it}=1) + P_\tau(C_{it}=0)\right].$$

Note that

$$\begin{aligned}
E_\tau(Y_{it} \mid D_i=0) &= E_\tau(Y_{it} \mid D_i=0, C_{it}=1)P_\tau(C_{it}=1 \mid D_i=0) + E_\tau(Y_{it} \mid D_i=0, C_{it}=0)P_\tau(C_{it}=0 \mid D_i=0) \\
&= E_\tau(Y_{it} \mid D_i=0, C_{it}=1)P_\tau(C_{it}=1 \mid D_i=0) \\
&= E_\tau(Y_{it} \mid C_{it}=1)P_\tau(C_{it}=1) \\
&= E_\tau(Y_{it}^* \mid C_{it}=1)P_\tau(C_{it}=1),
\end{aligned}$$

where the first equality follows from the law of iterated expectations and frame separability, the second equality follows from frame monotonicity, the third equality follows from frame exogeneity, and the fourth equality follows from the consistency principle. Analogously, it follows that

$$E_\tau(Y_{it} \mid D_i=1) = E_\tau(Y_{it}^* \mid C_{it}=1)P_\tau(C_{it}=1) + P_\tau(C_{it}=0).$$

Combining the previous two equation and the bound above deliver the desired result. □

Proof of Proposition 2. First, we calculate the fraction of partition-consistent investors at x as follows:

$$\begin{aligned}
E_\tau(C_{it}^\theta(x)) &= P_\tau\left(\{\theta_{it}(0) \le x \cap \theta_{it}(1) \le x\} \cup \{\theta_{it}(0) > x \cap \theta_{it}(1) > x\}\right) && (25) \\
&= P_\tau(\theta_{it}(1) \le x) + P_\tau(\theta_{it}(0) > x), && (26)
\end{aligned}$$

where the first line follows by definition and the second line follows from the properties of a probability measure and frame monotonicity. Next, we calculate the counter-CDF of the preferences of partition-consistent investors at x:

$$\begin{aligned}
P_\tau(\theta_{it}^* > x \mid C_{it}^\theta(x)=1) &= \frac{P_\tau(\theta_{it}^* > x \cap C_{it}^\theta(x)=1)}{P_\tau(C_{it}^\theta(x)=1)} && (27) \\
&= \frac{P_\tau(\theta_{it}(0) > x)}{P_\tau(\theta_{it}(1) \le x) + P_\tau(\theta_{it}(0) > x)}, && (28)
\end{aligned}$$

where the first line follows by definition and the second line follows from monotonicity and (26). Using these previous two results, we can now derive bounds on the CDF of preferences in the population:

$$\begin{aligned}
P_\tau(\theta_{it}^* \le x) &= P_\tau(\theta_{it}^* \le x \mid C_{it}^\theta(x)=1)P_\tau(C_{it}^\theta(x)=1) + P_\tau(\theta_{it}^* \le x \mid C_{it}^\theta(x)=0)P_\tau(C_{it}^\theta(x)=0) \\
&= P_\tau(\theta_{it}(1) \le x) + P_\tau(\theta_{it}^* \le x \mid C_{it}^\theta(x)=0) * P_\tau(C_{it}^\theta(x)=0) \\
&= P_\tau(\theta_{it}(1) \le x) + P_\tau(\theta_{it}^* \le x \mid C_{it}^\theta(x)=0) * \left[1 - P_\tau(\theta_{it}(1) \le x) - P_\tau(\theta_{it}(0) > x)\right]
\end{aligned}$$

$$\in [P_\tau(\theta_{it}(1) \leq x), P_\tau(\theta_{it}(1) \leq x) + 1 - P_\tau(\theta_{it}(1) \leq x) - P_\tau(\theta_{it}(0) > x)]$$
$$= [P_\tau(\theta_{it}(1) \leq x), P_\tau(\theta_{it}(0) \leq x)]$$

where the second equality follows from (28) and the third equality follows from (26). Combining the previous equation with the definition of $E_\tau(\cdot)$ delivers the desired result. □

Proof of Proposition 3. By the consistency principle,

$$E_\tau(Y_{it}^* \mid C_{it}^Y = 1) = E_\tau(Y_{it} \mid C_{it}^Y = 1).$$

By the law of iterated expectations,

$$E_\tau(Y_{it} \mid C_{it}^Y = 1) = E_\tau(Y_{it} \mid C_{it}^Y = 1, Y_{it} = D_i) P_\tau(Y_{it} = D_i \mid C_{it}^Y = 1) + E_\tau(Y_{it} \mid C_{it}^Y = 1, Y_{it} \neq D_i) P_\tau(Y_{it} \neq D_i \mid C_{it}^Y = 1).$$

Frame exogeneity implies the two expectations on the right-hand side of the previous equation are equal to $E_\tau(Y_{it} \mid C_{it}^Y = 1)$, which delivers the desired result. An identical argument follows for stock shares. □

Derivation of (6) and (7). These expressions follow from the following identity, which applies when W is binary:

$$cov(V, W) = E(VW) - E(V)E(W) = E(W)[E(V|W = 1) - E(V)].$$

□

Appendix C. Additional Results

Figure A1. Preferences vs. Frictions: Stock Share

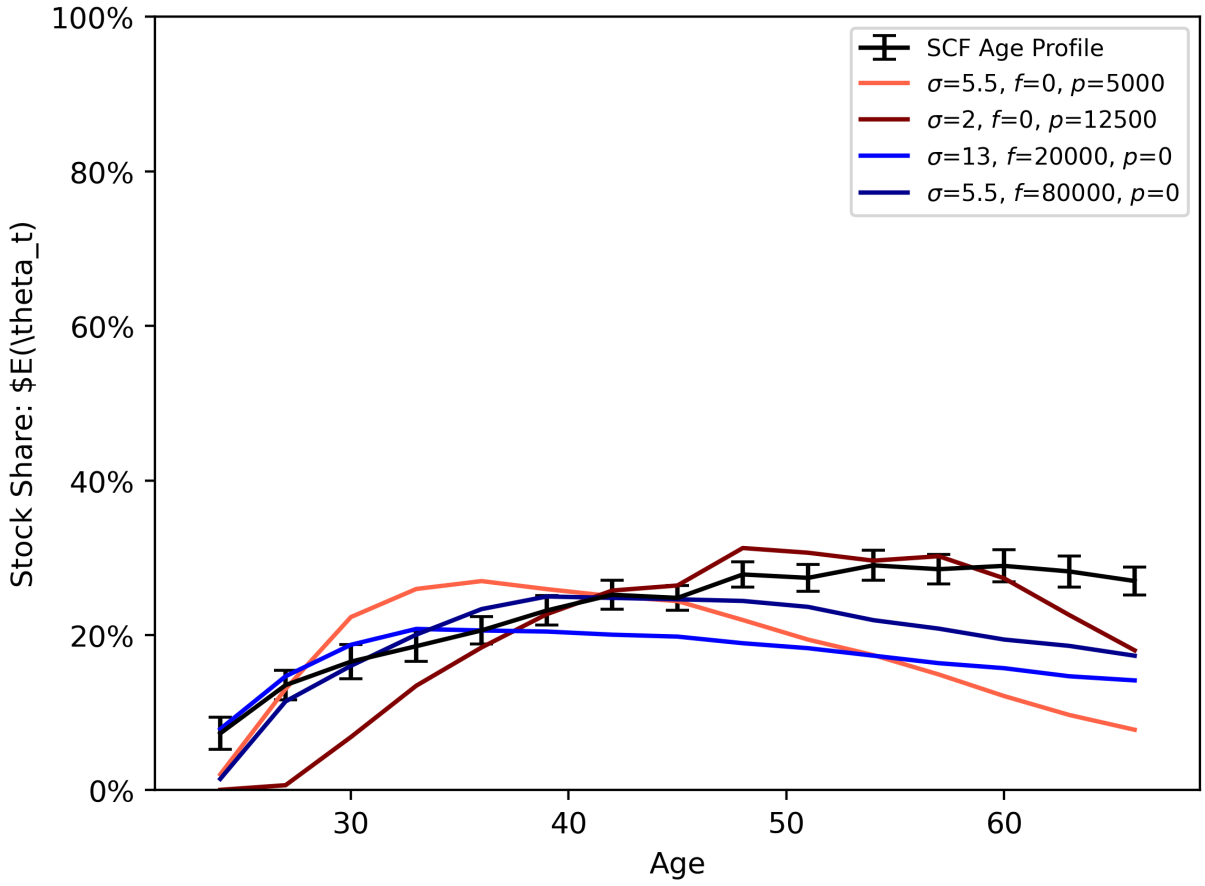


Figure A2. Distribution of Treatment and Control Groups by Year: Opt-In to TDF Sample

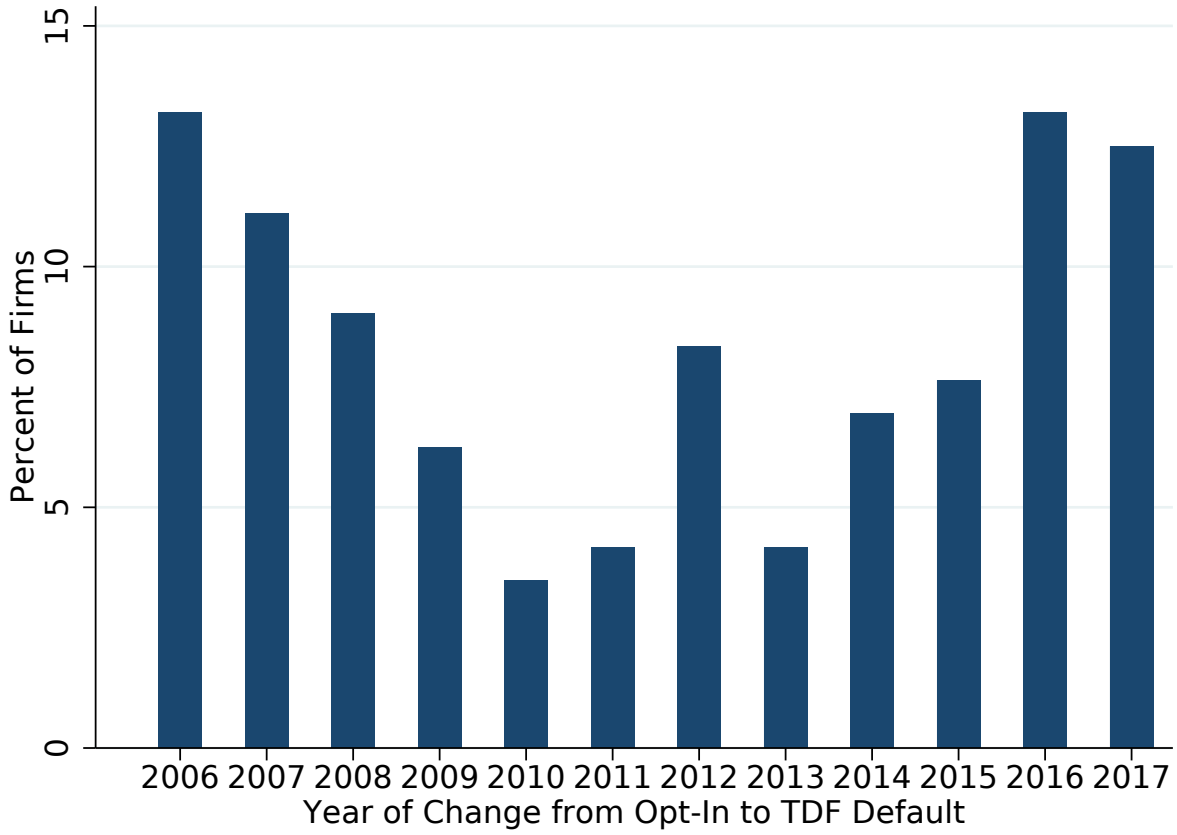
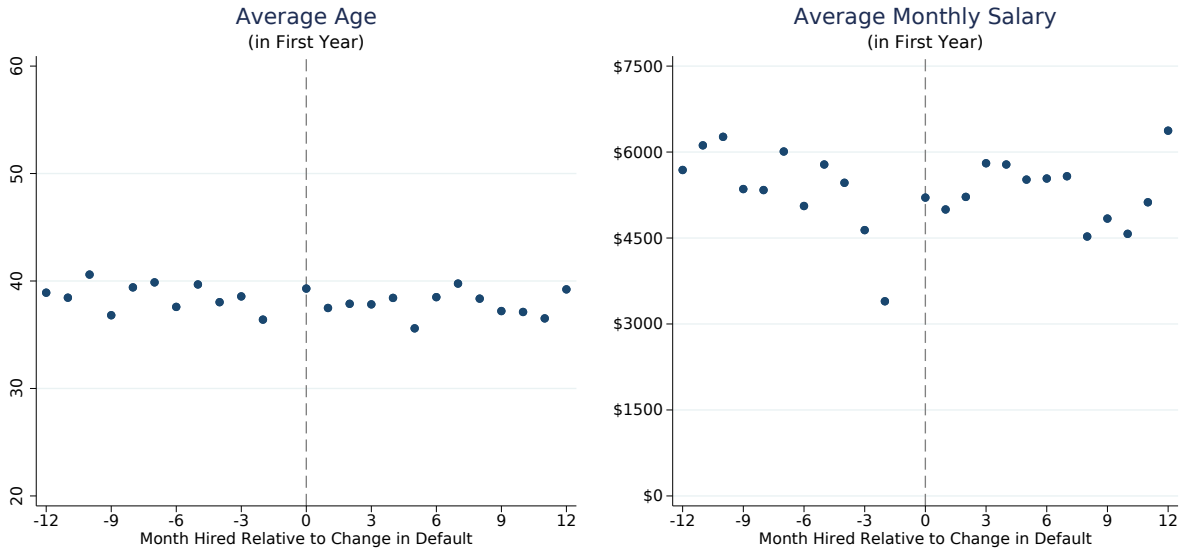


Figure A3. Balance Checks

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

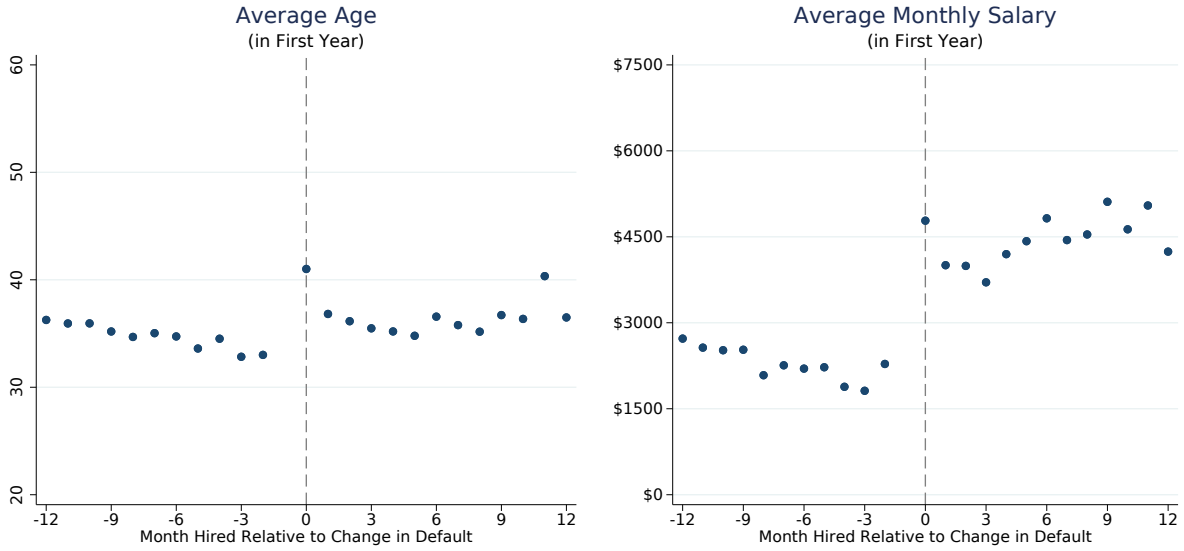
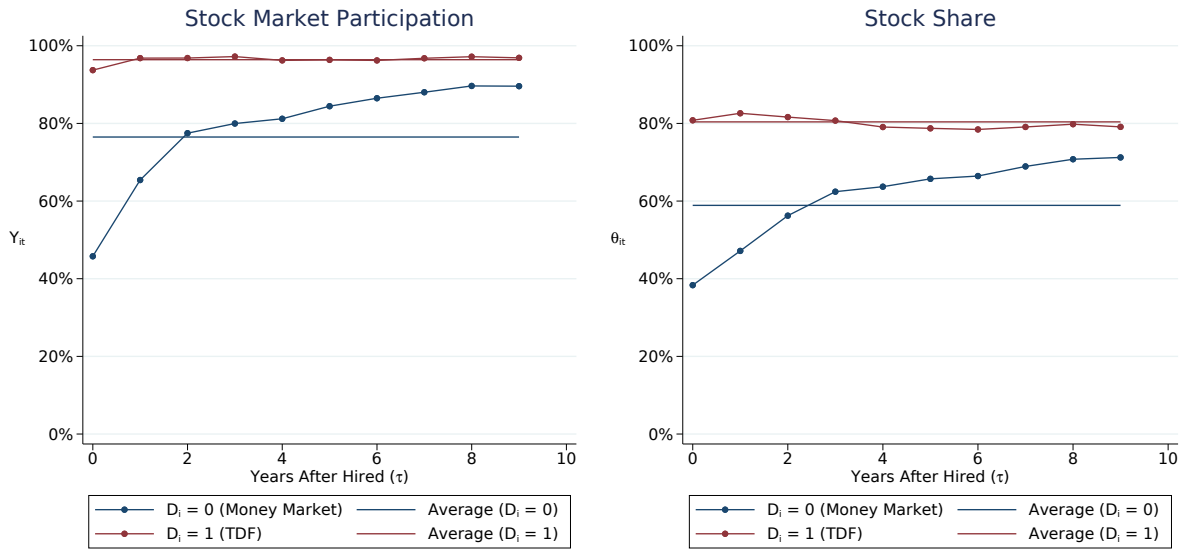


Figure A4. Observed Portfolio Choice Response

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

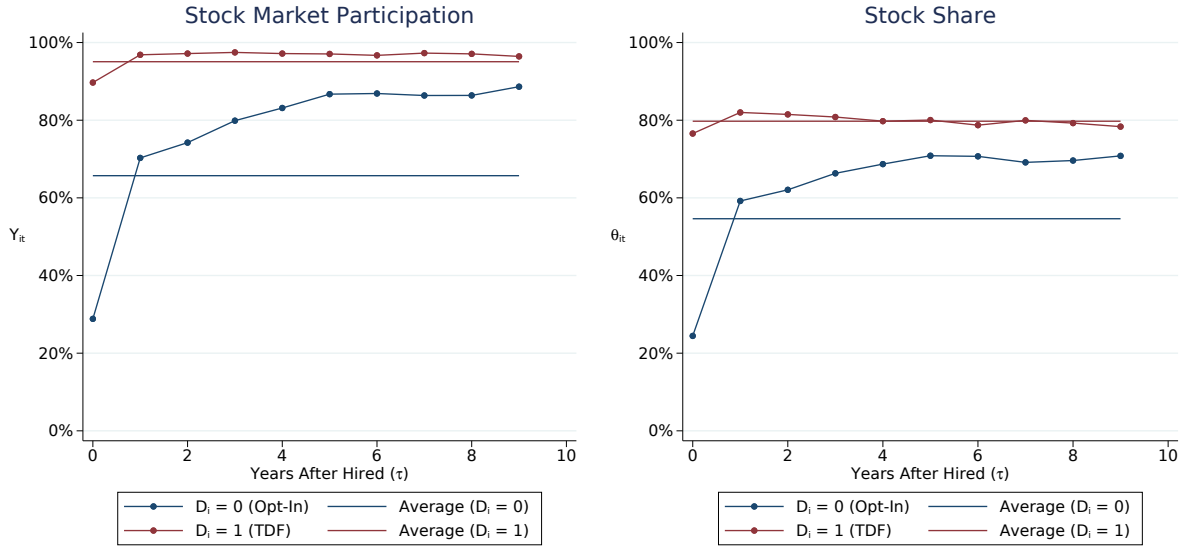
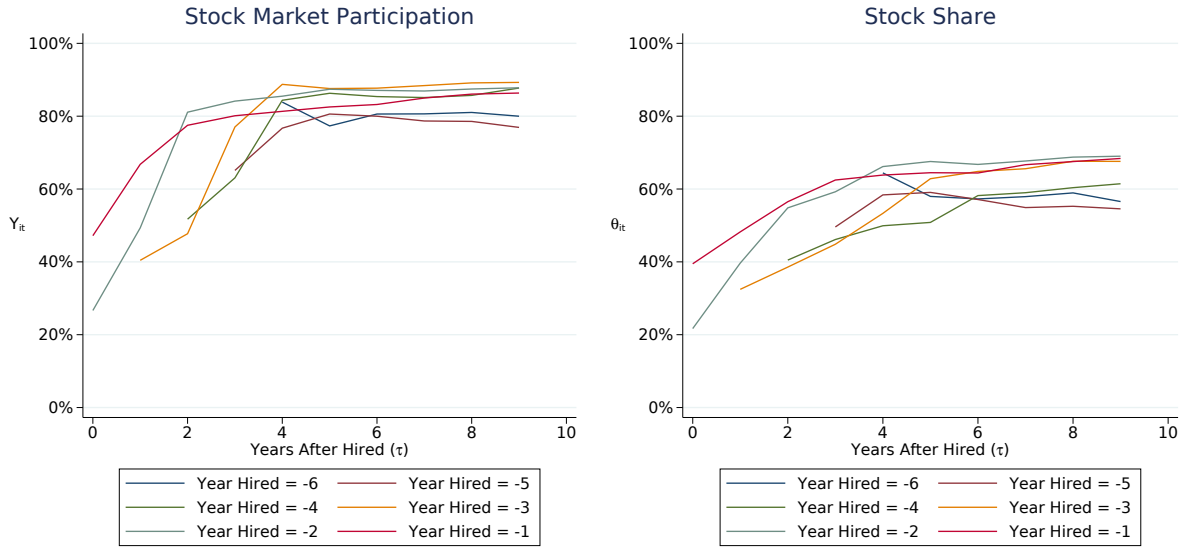


Figure A5. Robustness of Portfolio Choice Response: Money Market to TDF Sample

Panel A: Peer Effects

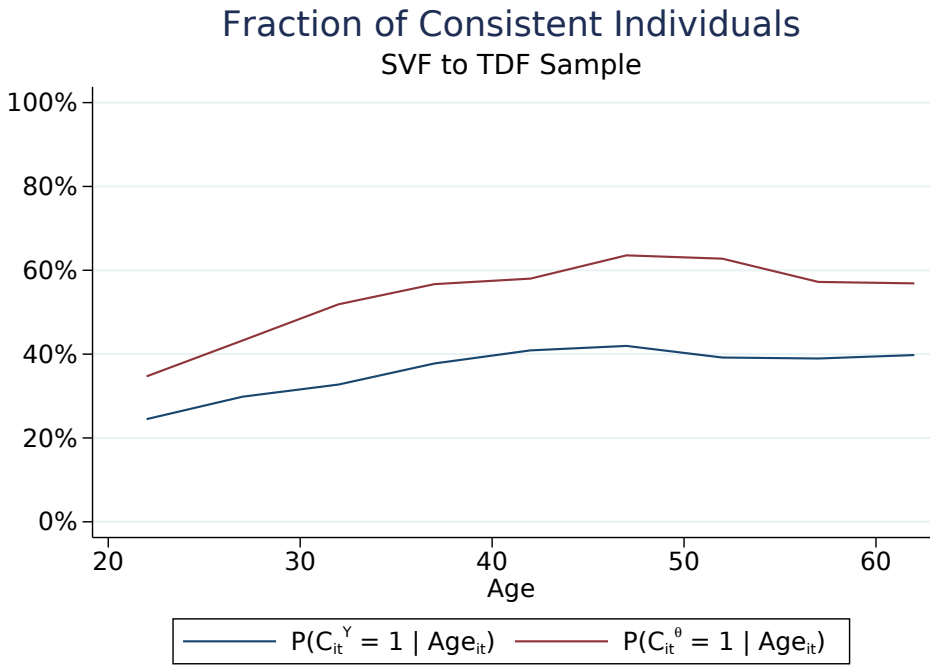


Panel B: Compositional Change



Figure A6. Fraction of Consistent Investors by Age

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

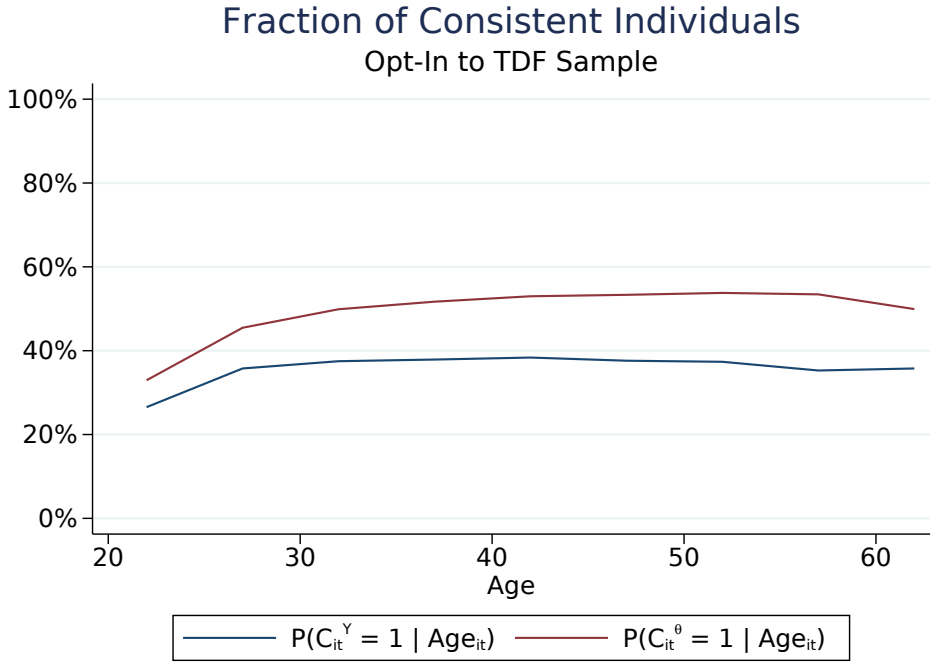
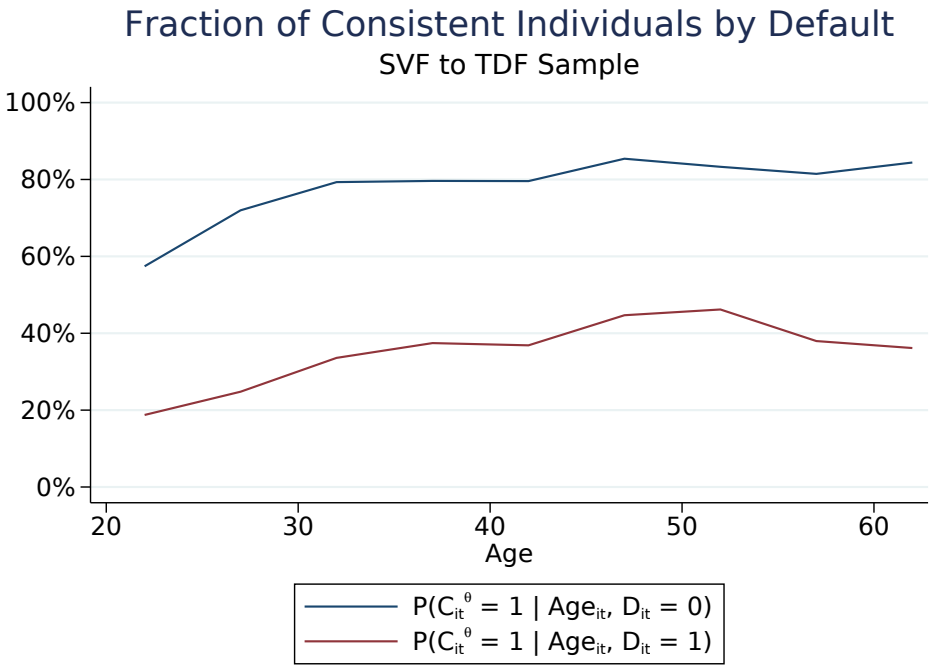


Figure A7. Fraction of Consistent Investors by Age and Default

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

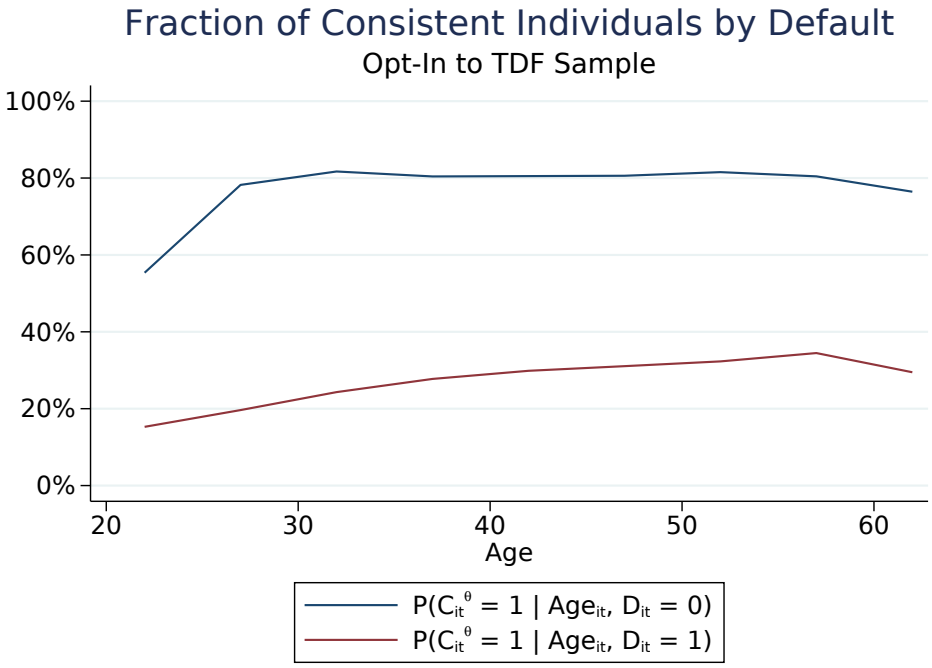
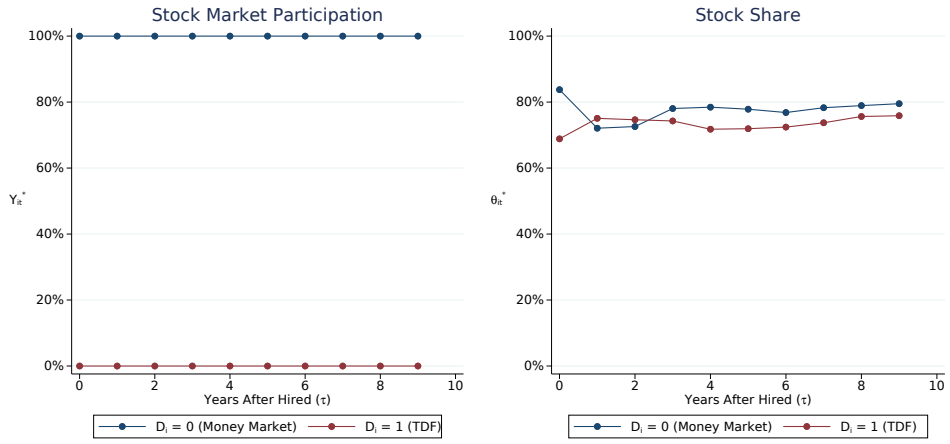


Figure A8. Preferences of Consistent Investors by Default

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

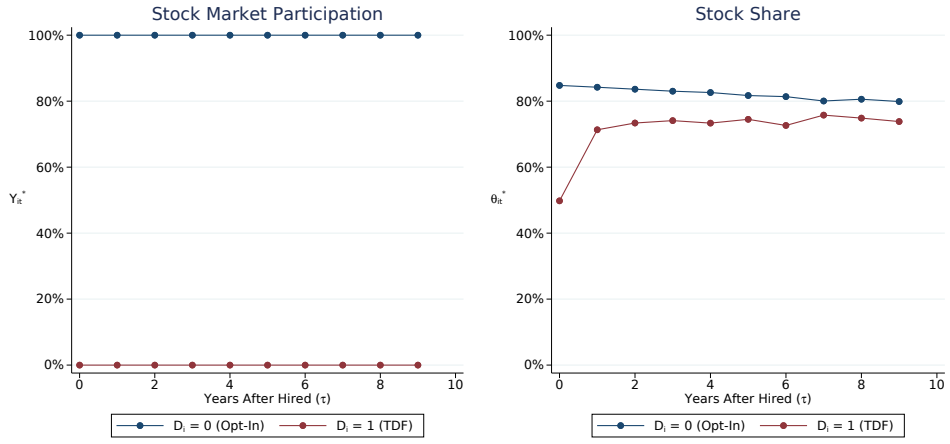
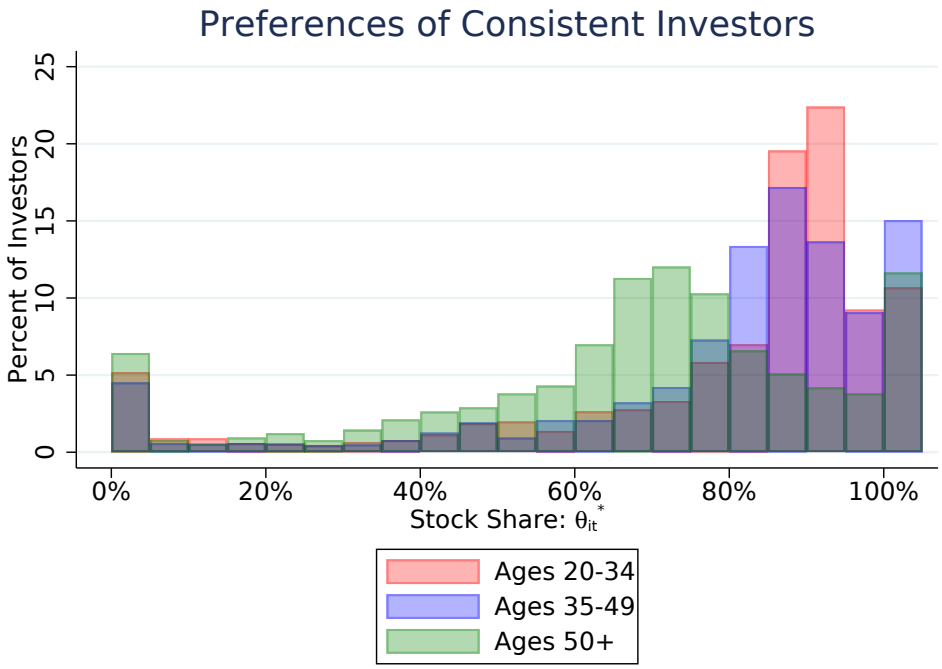


Figure A9. Preference Heterogeneity among Consistent Investors

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

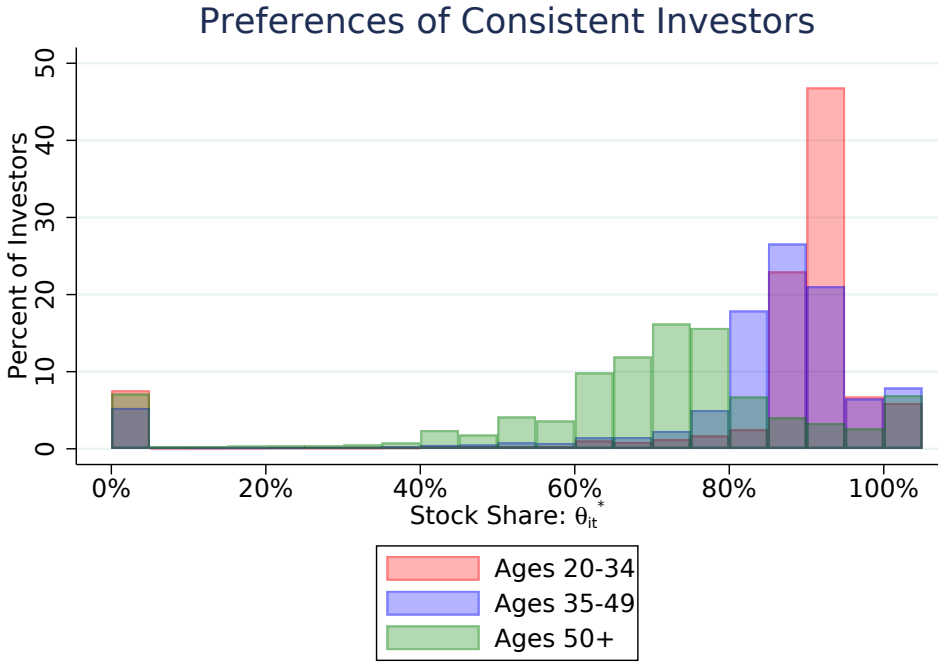
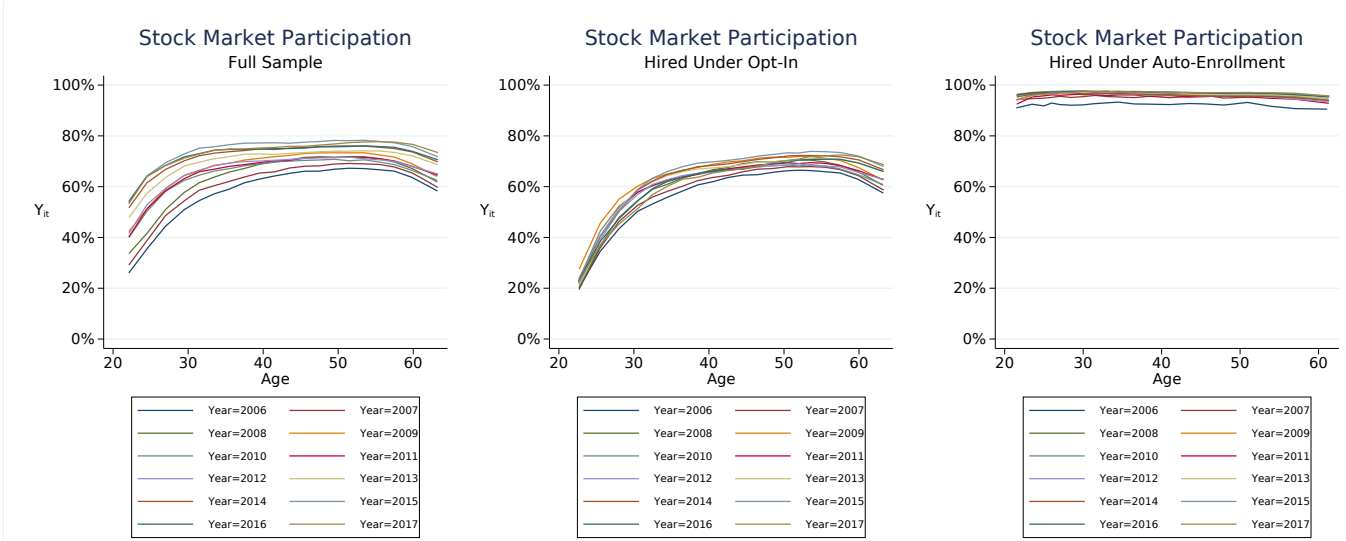


Figure A10. Year Effects in Observed Choices

Panel A: Participation



Panel B: Stock Share

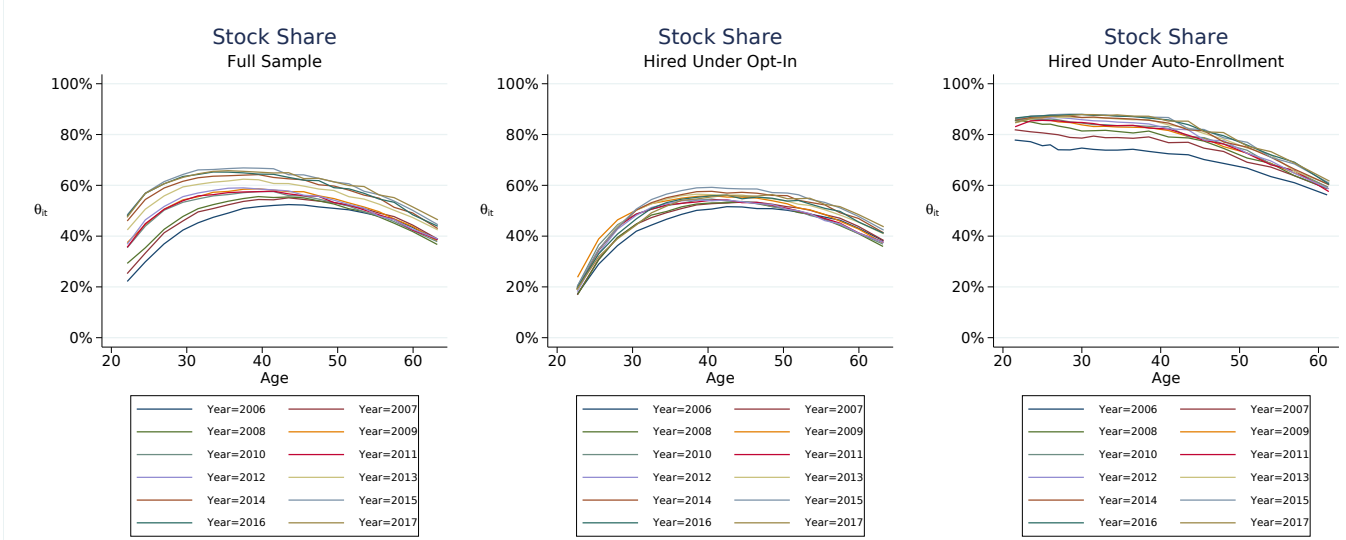
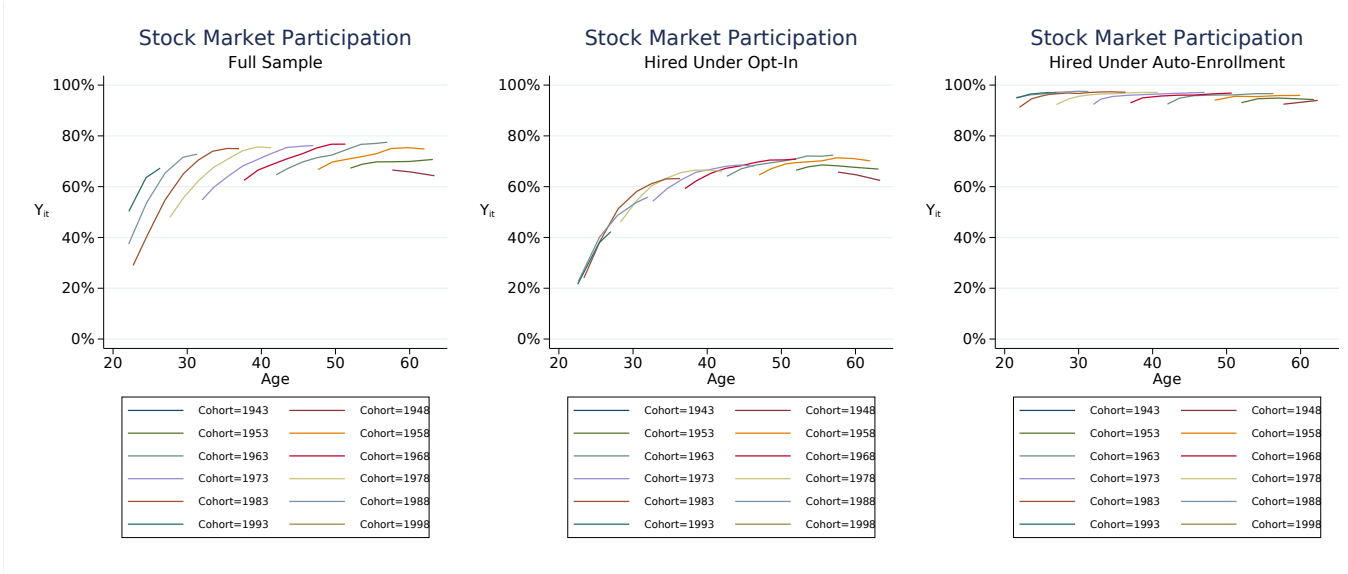


Figure A11. Cohort Effects in Observed Choices

Panel A: Participation



Panel B: Stock Share

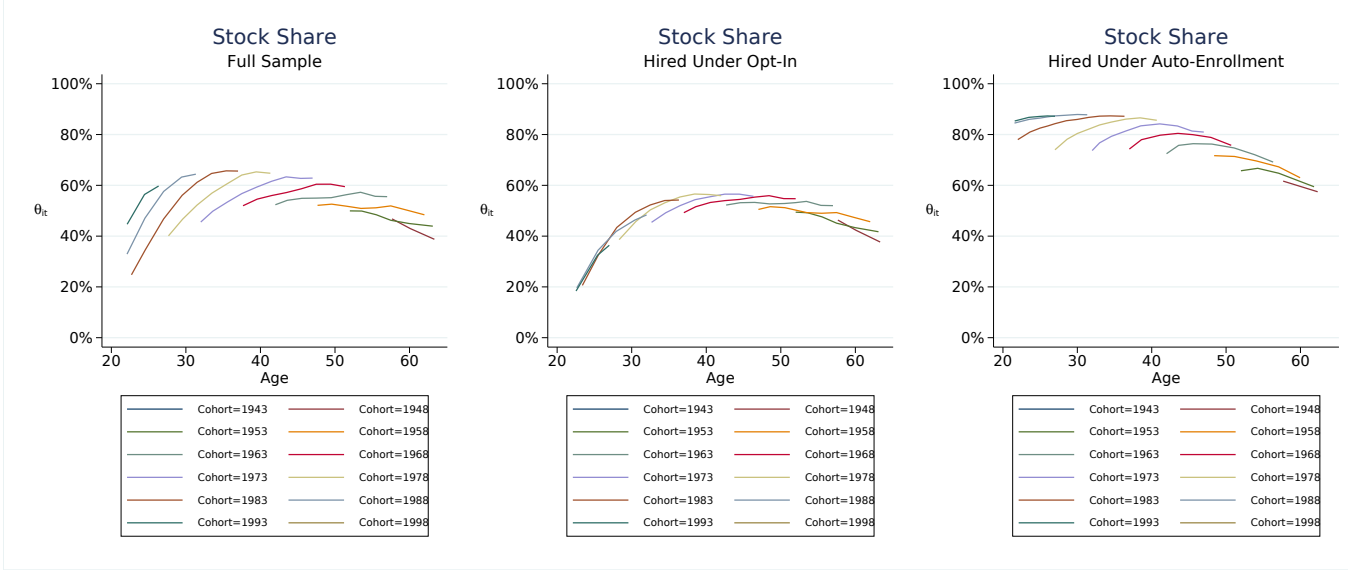
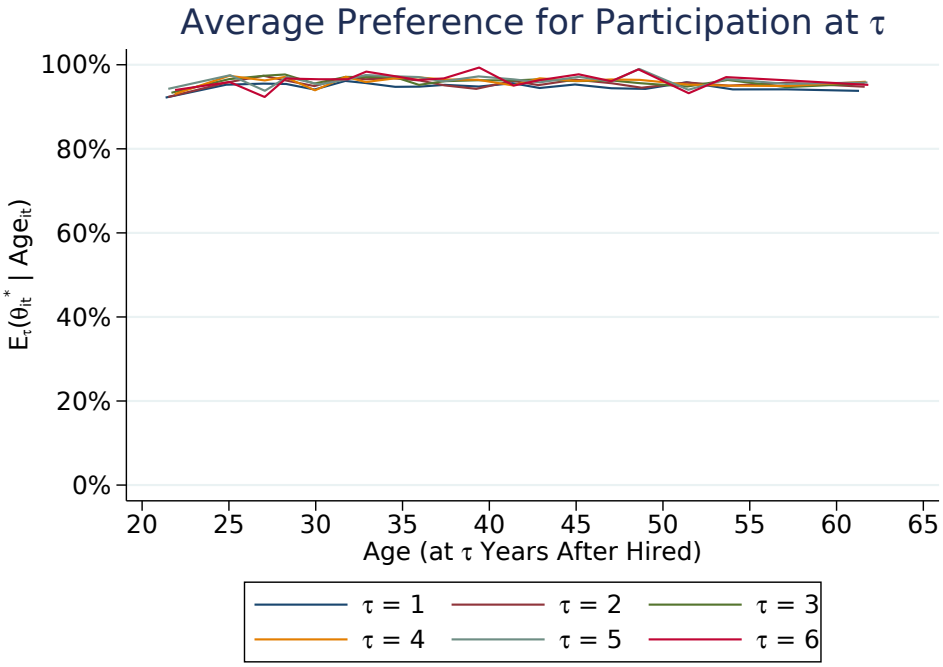


Figure A12. Robustness of Preferences over the Life Cycle: Opt-In to TDF Sample

Panel A: Participation



Panel B: Stock Share

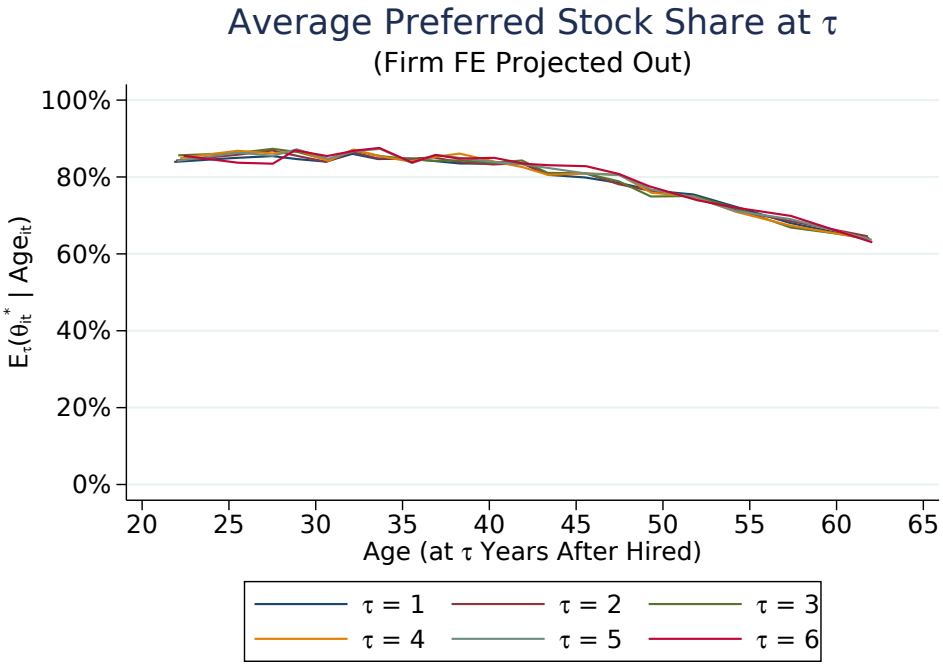
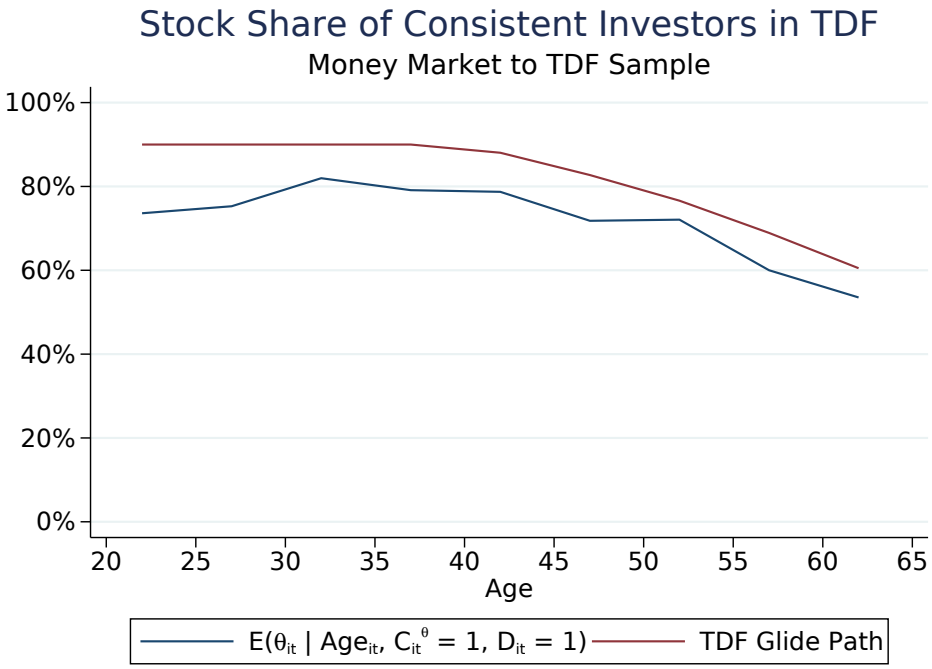


Figure A13. Life Cycle Preferences of Consistent Investors Defaulted into TDF

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

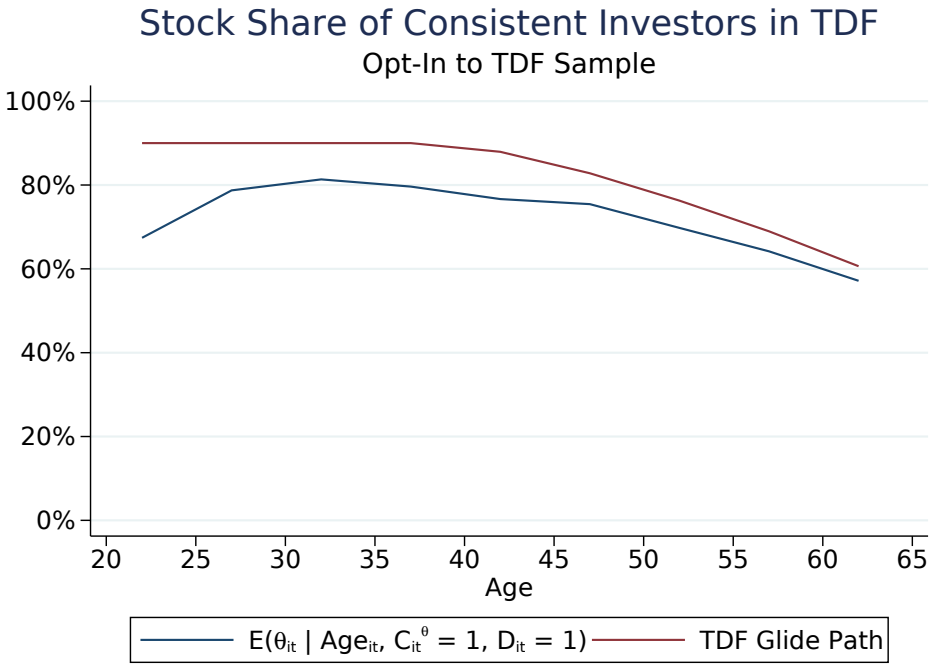
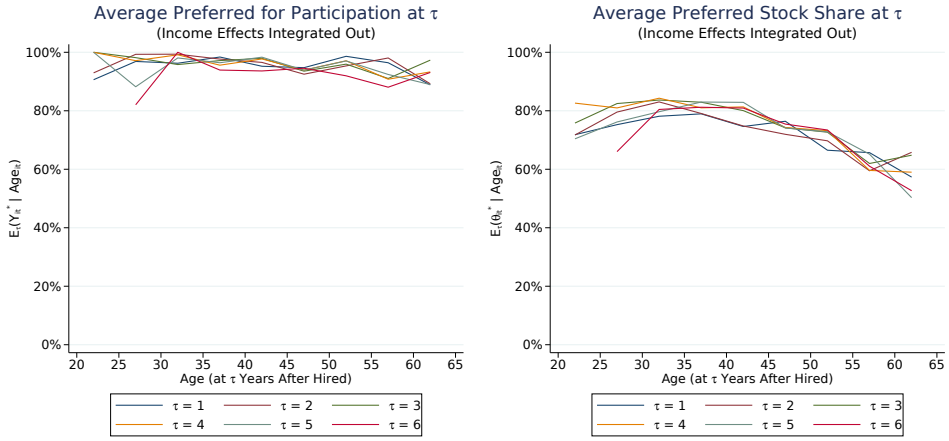


Figure A14. Estimated Preferences Under Weaker Identifying Assumption

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample

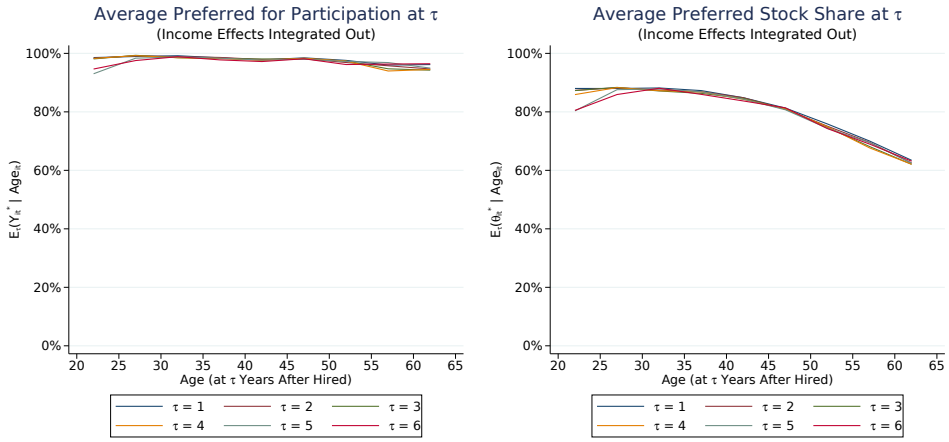
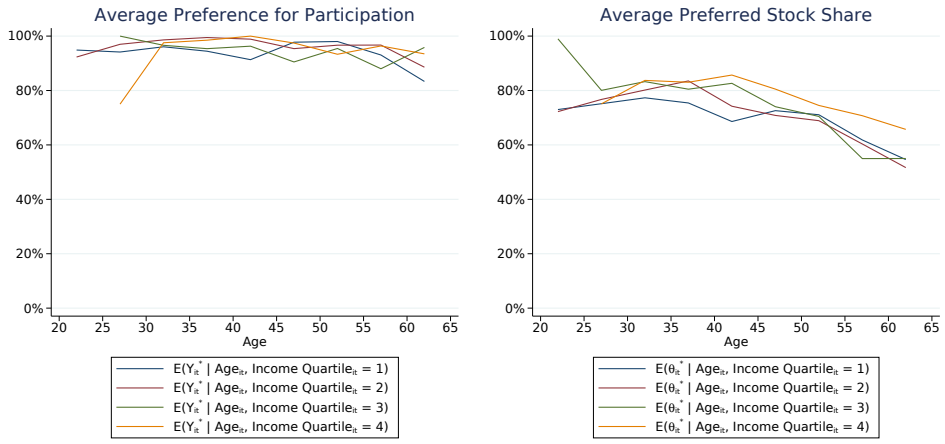
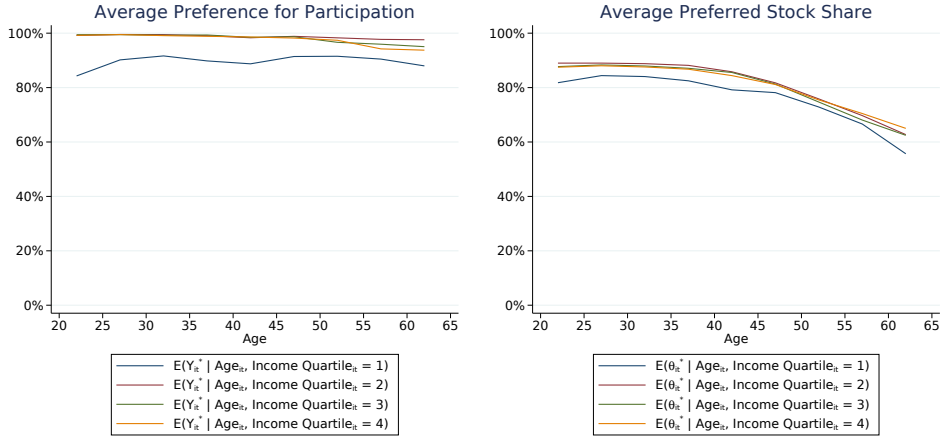


Figure A15. Preferences over the Life Cycle by Income Quartiles

Panel A: Money Market to TDF Sample



Panel B: Opt-In to TDF Sample



Appendix D. Life Cycle Model Parameters

Preference Parameters	
δ	Discount factor
σ	Relative Risk Aversion
k_s	Contribution adjustment cost
k_θ	Portfolio adjustment cost
Utility	
$u_t(\,)$	Utility function
$V(\,)$	Value function
State variables	
X_t	Vector of all state variables
a	Age
emp_t	Employment status
ten_t	Tenure
e	Employer DC plan type
η	Labor productivity
ae	Average lifetime earnings
L	Liquid assets
A	DC wealth stock
s_d	Default contribution rate
Θ_d	Default allocation existing funds
θ_d	Default allocation new contributions
$F(\,)$	Joint distribution of $\eta, emp, e, \varepsilon$
Choices	
c	Consumption
s^{dc}	DC contribution rate
s^l	Savings in liquid assets
Θ	Asset allocation for existing funds
θ	Asset allocation for new contributions
Defined Contribution Account	
$F_{\overline{\theta}}(\,), F_{\overline{s}}(\,)$	Distributions of DC account defaults
$\mathcal{M}_e(\,)$	Employer DC matching function
$fixed_e$	Fixed employer contribution
$match_e$	Employer matching rate
cap_e	Threshold on employer matching
$\Upsilon_e(\,)$	Vesting risk-adjustment

Assets	
$R^{liq}(\,)$	Rate of return on liquid assets
R_f	Risk-free rate
$1 + r^{cc}$	Interest rate on unsecured debt
$\underline{L}$	Borrowing limit
$R^S()$	Return on risky assets
μ_s	Log-risk premium
σ_s	S.D. of Log risky asset returns
Labor Market	
$\pi^{JJ}(\,)$	Job-to-job transition probability
$\pi^{EU}(\,)$	Unemployment transition probability
$\pi^{UE}(\,)$	Out-of-unemployment transition probability
$\{\delta_i\}_{i=0}^3$	Deterministic component of earnings
w	Labor earnings
ρ	Autocorrelation in earnings shocks
ξ	Earnings innovation if continuously employed
$\sigma^2_{\xi_0}$	Variance of the first earnings innovation
σ^2_ξ	Variance of subsequent innovations
ξ^{JJ}	Earnings innovation after job-to-job transition
μ^{JJ}	Avg. wage gain after a job-to-job transition
ξ^U	Earnings innovation out of unemployment
μ^{UE}	Avg. wage loss out of unemployment
ι	Measurement error in earnings
σ^2_ι	Variance of measurement error
ϕ	Earnings innovation plus measurement error
Demographics	
$\overline{T}$	Maximum years of life
T^w	Number of working years
m_t	Mortality risk
n_t	Equivalence scale
Tax and benefit system	
$tax_i(\,)$	Tax on income
τ^k	Tax on capital returns
$limit_a$	Tax limit on DC contributions
pen_a	Tax penalty of early DC withdrawals
$ui(\,)$	Unemployment insurance benefit
$ss(\,)$	Public pension income

Appendix E. First-Stage Estimation Details

E.1 Demographics

Survival probabilities. Survival probabilities for each age are calibrated to the U.S. Social Security 2015 Actuarial Life Tables.

Equivalence scale. Changes in household composition over the life cycle are captured by an equivalence scale in the utility function. We use the equivalence scale by age estimated by Lusardi et al. (2017). Using PSID data from 1984 to 2005, Lusardi et al. (2017) estimate $z(j_t,k_t) = (j_t + 0.7k_t)^{0.75}$ where j_t and k_t are, respectively, the average number of adults and children (under 18 years old) in a household with a head of age t. They normalize this measure by $z(2,1)$—the composition of a household with 2 adults and 1 child—to get the equivalence scale at age t equal to $n_t = \frac{z(j_t,k_t)}{z(2,1)}$. To estimate n_t we use publicly available replication files from Lusardi et al. (2017) and aggregate the data across education groups.

E.2 Taxes and Benefit System

Income taxation. Taxable income is defined as the sum of labor earnings, social security and unemployment benefits, DC withdrawals, less contributions to the DC account:

$$y_t^{tax} = \begin{cases} w_t - s_t^{dc} * w_t & if\ emp_t \in \{E, JJ\} \\ ui_t + s_t^{dc} * w_t & if\ emp_t = U \\ ss\left(ae_{T^w}\right) + s_t^{dc} * w_t & if\ emp_t = Ret \end{cases}$$

Investors' income tax liability is calculated according to the federal income tax schedule of 2006 (the first year of data and the base year for the calibration) for an investor filling as single and claiming the standard deduction. The tax formula has 5 annual income brackets $\{\tilde{\kappa}_i^\tau\}_{i=1}^5 = \{\$5,150; \$7,550; \$30,650; \$74,200; \$154,800\}$.[26] Quarterly tax brackets are defined as: $\kappa_i^\tau = \frac{1}{4}\tilde{\kappa}_i^\tau$. The quarterly income tax liability is given in the following equation, which we aggregate to an annual frequency by multiplying by four.

[26] Note that the first bracket correspond to the standard deduction amount in 2006.

$$tax_t^i = \begin{cases} 0 & if\ y^{tax} \leq \kappa_1^\tau \\ 0.10\left(y^{tax} - \kappa_1^\tau\right) & if\ \kappa_2^\tau \geq y^{tax} > \kappa_1^\tau \\ 0.10\left(\kappa_2^\tau - \kappa_1^\tau\right) + 0.15\left(y^{tax} - \kappa_2^\tau\right) & if\ \kappa_3^\tau \geq y^{tax} > \kappa_2^\tau \\ 0.10\left(\kappa_2^\tau - \kappa_1^\tau\right) + 0.15\left(\kappa_3^\tau - \kappa_2^\tau\right) + 0.25\left(y^{tax} - \kappa_3^\tau\right) & if\ \kappa_4^\tau \geq y^{tax} > \kappa_3^\tau \\ 0.10\left(\kappa_2^\tau - \kappa_1^\tau\right) + 0.15\left(\kappa_3^\tau - \kappa_2^\tau\right) + 0.25\left(\kappa_4^\tau - \kappa_3^\tau\right) + 0.28\left(y^{tax} - \kappa_4^\tau\right) & if\ \kappa_5^\tau \geq y^{tax} > \kappa_4^\tau \\ 0.10\left(\kappa_2^\tau - \kappa_1^\tau\right) + 0.15\left(\kappa_3^\tau - \kappa_2^\tau\right) + 0.25\left(\kappa_4^\tau - \kappa_3^\tau\right) + 0.28\left(\kappa_5^\tau - \kappa_4^\tau\right) + 0.33\left(y^{tax} - \kappa_5^\tau\right) & if\ y^{tax} > \kappa_5^\tau \end{cases}$$

Public pension. The amount of public pension benefit (ss) is computed according the 2006 Social Security formula at the full retirement age, with an income floor guaranteed by the Supplemental Security Income program (with a monthly benefit si = \$603). Annual public pension benefits are equal to:

$$ss\left(ae_{T^w}\right) = 4 \times 3 \times \max\left\{si; \tilde{ss}\left(ae_{T^w}\right)\right\}$$

where $\tilde{ss}$, the monthly social security benefit, is increasing in average lifetime earnings ae_{T^w} up to a maximum monthly benefit:

$$\tilde{ss} = \begin{cases} 0.90 \times \frac{1}{3} ae_{T^w} & if\ \frac{1}{3} ae_{T^w} \leq \$656 \\ 0.90 \times \$656 + 0.32 \times \left(\frac{1}{3} ae_{T^w} - \$656\right) & if\ \$3,955 > \frac{1}{3} ae_{T^w} > \$655 \\ \min\left\{0.90 \times \$656 + 0.32 \times \$3,299 + \left(0.15 \times \frac{1}{3} ae_{T^w} - \$3,299\right); \$2,053\right\} & if\ \frac{1}{3} ae_{T^w} > \$3,955 \end{cases}$$

Unemployment benefits. Unemployment insurance provides a constant replacement rate ω of labor earnings implied by the labor productivity level in the last period of employment. Labor productivity η_t stays constant during an unemployment spell. We set $\omega = 0.40$, which is the average replacement rate across all U.S. states (U.S. Department of Labor, 2018). For simplicity, we assume that the employer contribution portion of an early withdrawal is always equal to the employer match rate. This simplifying assumption is valid assuming participants contribute below the matching threshold and contributions are fully vested. Adjusted unemployment benefits for an investor unemployed since period $t - x$ are given by:

$$ui_t = \max\left\{0; \omega w_t\left(\eta_{t-x}\right) - s_t^{dc} * w_t\right\}$$

E.3 Assets

Asset returns and limits. The properties for financial assets are described in the main text. In the baseline estimation described in the main text, we set $\underline{L}_t = 0$ for all ages.

Asset taxation. In line with IRS rules for 2006, the maximum contribution limit for tax-deferred retirement contributions ($limit_a$) is set equal to \$15,000 annually for investors younger than 50 years old and \$20,000 after that in 2006

dollars. The tax penalty for early DC withdrawals (pen_t) is equal to 10% before age 55 and to zero afterwards[27].Capital income is taxed at a rate τ^k of 15% which implies that the (net of tax) annual returns on liquid wealth is 2.55%.

Vesting schedule. An employee who separates from her employer before the end of the vesting period may lose part (or all) of the employer matching contribution. A vesting schedule $vst_e(ten)$ determines the percentage of employer contributions that an employee keeps if she separates at a given tenure level. Modeling the vesting schedule explicitly would introduce an additional continuous state variable to the dynamic problem: the amount of non-vested of DC wealth. Instead, we adjust employer contributions by a factor $\Upsilon_e(t,ten)$ proportional to the risk of losing unvested employer contributions. The adjustment factor $\Upsilon_e(t,ten)$ is given in equation (29). It depends on both the cumulative job-separation probability and the vesting schedule. It is smaller than one and increasing in tenure before the end of the vesting period, and equal to one afterward. Importantly, this specification captures the fact that vesting matters more for employees who—based on their age and tenure—are more likely to separate from their employer.

$$\Upsilon_e(t,ten) = 1 - \sum_{j=0}^{T^R-t} \left(\prod_{k=1}^{j-1} \left(1 - \pi^{EU}_{t+k,ten+k} - \pi^{JJ}_{t+k,ten+k}\right) \right) \left(\pi^{EU}_{t+j,ten+j} + \pi^{JJ}_{t+j,ten+j}\right) \left(1 - vst_e(ten+j)\right) \tag{29}$$

On average, in the estimation sample of 34 U.S. 401(k) plans, 52% of matching contribution are vested immediately and this share increases over tenure. The average vested share reaches 70% by the end of the second year of tenure. We assume that all matching contributions are fully vested starting from the 9th quarter of tenure.

E.4 Labor Market Parameters

Our estimation of the labor market parameters follows Choukhmane (2021). In this sub-section, we provide a brief description of this estimation, but refer the reader to Choukhmane (2021) for additional details.

Data. We use the Survey of Income and Programs and Participation (SIPP) to estimate of the wage earnings process and labor market transitions probabilities. We use the 1996 panel of the SIPP which contains data from December 1995 to February 2000 and aggregate the data at quarterly frequency. We focus on an investor's primary job (defined as the job where he worked the most hours). We restrict the sample to investors aged 22 to 65 years old, and exclude full-time students and business owners. We assign employment status based on investors' responses in the first week of each quarter. An investor is classified as employed if she reports having a job. We record a job-to-job transition if the identity of an investor's employer is different in two successive quarters. We record a job separation if an investor is employed in the beginning of a quarter, and not employed in the beginning of the subsequent quarter. Job separations include early retirement decisions, before the age of 65.

Earnings process. We estimate the labor earnings process for workers staying in the same job using a standard two-step minimum distance approach similar to Guvenen (2009) and Low, Meghir, and Pistaferri (2010). The empirical income process is given in equation (30), which is the empirical counterpart of the model earning process in equation

[27]In the model, early withdrawals are only allowed in periods of unemployment. The tax code allows penalty-free 401(k) hardship withdrawals for unemployed people older than 55, which is earlier than the normal 59$^1/_2$ eligibility age for penalty-free withdrawals.

(10) with one additional term: serially uncorrelated measurement error $\iota_{i,t} \sim N\left(0, \sigma_\iota^2\right)$.

$$\ln w_{i,t} = \delta_0 + \delta_1 a_{i,t} + \delta_2 a_{i,t}^2 + \delta_3 a_{i,t}^3 + \underbrace{\phi_{i,t}}_{\eta_{i,t}+\iota_{i,t}} \tag{30}$$

The estimation has two steps. In the first step, We estimate the parameters of the deterministic component of earnings $\left(\{\delta_j\}_{j=0}^3\right)$—a cubic in age. In the second step, We use the residual from regression (30) to estimate the five parameters governing the stochastic component of earnings: the coefficient of autocorrelation in earnings shocks (ρ), the variances of the first earnings innovation ($\sigma_{\xi_0}^2$), the variance of subsequent innovations (σ_ξ^2), and the variance of measurement error (σ_ι^2). We estimate these five parameters by minimizing the distance between the empirical variance-covariance matrix of earnings residuals and its theoretical counterpart implied by the statistical model. The resulting estimates are provided in Table A1.

Table A1. Earnings Process Estimates

This table shows quarterly earnings process estimated using a two-step minimum distance estimator on a panel of workers continuously employed in the same job. Data source: U.S. Survey of Income and Program Participation.

Age component				**Stochastic component of earnings**			
δ_0	δ_1	δ_2	δ_3	ρ	$\sigma_{\xi_0}^2$	σ_ξ^2	σ_ι^2
1.632	0.094	−0.00107	$3.9e^{-7}$	0.974	0.184	0.0125	0.10

Earnings after a transition. We estimate the median change in log salary following a job-to-job transition (μ^{JJ}) to be equal to 0.048. We estimate that job transitions following a period of unemployment are associated with a loss in earnings. We estimate the median change in log salary relative to the last salary prior to unemployment (μ^{UE}) to be equal to −0.078.

Numeraire. The average net compensation per worker in the U.S. was around \$37,078 in 2006 (from the Social Security Administration national average wage index). This is also almost equal to the median annual salary in the estimation sample (\$37,998 in 2006 dollars). We thus calibrate annual earnings to this numeraire.

Labor transition probabilities. We use SIPP micro-data to estimate quarterly job-to-job (π^{JJ}) and job to non-employment (π^{EU}) transition probabilities by age and tenure and job finding rates (π^{UE}) by age. The initial unemployment rate is set equal to 22%, which is the share not employed at age 22 in SIPP. The probability that an employed investor switches to another job (given in equation (31)) or moves to non-employment (given in equation (32)) is the sum of an age component (i.e. a fifth-order polynomial in age) and a tenure component (a set of dummies for employees in their first 9 quarters of tenure):

$$\pi^{JJ}(a,ten) = \sum_{k=1}^{5} \alpha_k^{JJ} a^k + \sum_{j=1}^{9} \iota_k^{JJ} 1\{(ten = j)\} \tag{31}$$

$$\pi^{EU}(a,ten) = \sum_{k=1}^{5} \alpha_k^{EU} a^k + \sum_{j=1}^{9} \iota_k^{UE} 1\{(ten = j)\} \tag{32}$$

The probability that an unemployed investor finds a job, given in equation (33), is defined as a fifth-order polynomial

in age.

$$\pi^{UE}(a) = \sum_{k=1}^{5} \alpha_k^{EU} a^k \tag{33}$$

We estimate equations (31), (32), and (33) using a linear probability regression. Estimates for the age component of labor market transitions are reported in Figure A16. Estimates for the tenure component are reported in Figure A17.

Figure A16. Age Component of Quarterly Labor Market Transitions

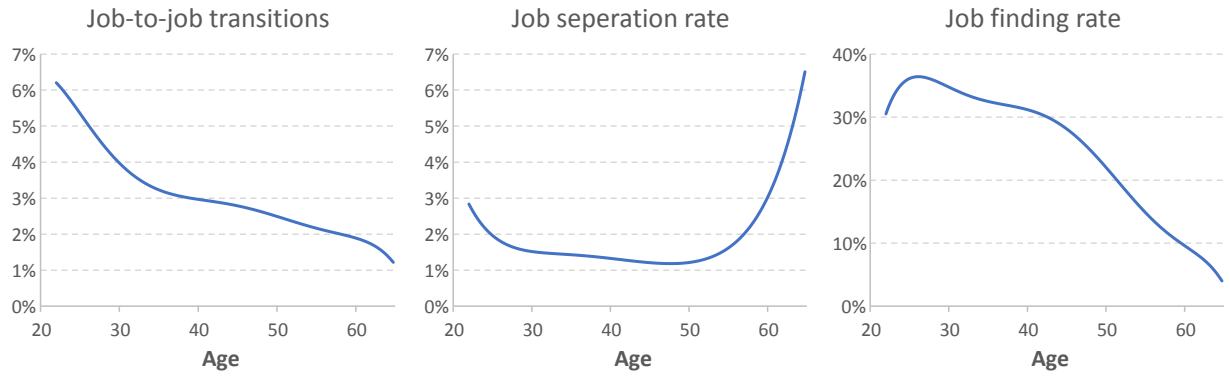


Figure A17. Tenure Component of Quarterly Labor Market Transitions

