

NBER WORKING PAPER SERIES

INCORPORATING EMPLOYEE HETEROGENEITY INTO DEFAULT RULES FOR
RETIREMENT PLAN SELECTION

Gopi Shah Goda
Colleen Flaherty Manchester

Working Paper 16099
<http://www.nber.org/papers/w16099>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2010

We would like to thank James Robertson for excellent research assistance, and Kristine Brown, Maria Fitzpatrick, Caroline Hoxby, Damon Jones, Ron Laschever, Brigitte Madrian, Kevin Mumford, Anita Alves Pena, Greg Rosston, and John Shoven for helpful comments as well as seminar participants at Stanford University and the University of Minnesota. This research was supported by the Center for Retirement Research at Boston College pursuant to a grant from the U.S. Social Security Administration funded as part of the Retirement Research Consortium through the Steven H. Sandell 2009-2010 Grant Program for Junior Scholars in Retirement Research. The opinions and conclusions are solely those of the authors and should not be construed as representing the opinions or policy of the Social Security Administration, any agency of the Federal Government, or the Center for Retirement Research at Boston College. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2010 by Gopi Shah Goda and Colleen Flaherty Manchester. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Incorporating Employee Heterogeneity into Default Rules for Retirement Plan Selection
Gopi Shah Goda and Colleen Flaherty Manchester
NBER Working Paper No. 16099
June 2010
JEL No. H8,J26

ABSTRACT

This paper examines the effect of incorporating individual-level heterogeneity into default rules for retirement plan selection. We use data from a large employer that transitioned from a defined benefit (DB) plan to a defined contribution (DC) plan, offering existing employees a one-time opportunity to make an irrevocable choice between plans. Employees who did not make a choice were defaulted to switch to the DC plan if under age 45 or remain in the DB plan if age 45 or older. Using a regression discontinuity framework, we estimate that the default increased the probability of enrolling in one plan over the other by 60 percentage points. We develop a framework to solve for the optimal age-based default rule analytically and use our results to empirically evaluate the optimal age-based default rule for the firm in our setting. Our simulations show that for a broad range of levels of risk aversion, allowing the default for the choice between pension plans to vary by age can substantially improve outcomes relative to a uniform default policy. Our results suggest that considerable welfare gains are possible in our model by varying defaults by observable characteristics.

Gopi Shah Goda
Stanford University
SIEPR
366 Galvez St.
Stanford, CA 94305
and NBER
gopi@stanford.edu

Colleen Flaherty Manchester
Human Resources and Industrial Relations
University of Minnesota
Room 3-300R
321 - 19th Avenue South
Minneapolis, MN 55455
cmanch@umn.edu

1 Introduction

The impact of default rules, which specify an individual's outcome when no choice is made, has been well-documented over the last decade. From dramatically increasing participation in 401(k) plans (e.g., Madrian and Shea 2001) to influencing organ donor status (e.g., Johnson and Goldstein 2003), default rules have a large effect on outcomes despite holding the set of choices constant. They have attracted widespread attention because their influential effects challenge the standard economic assumption of rational decision-making (DellaVigna 2009).

In this paper, we examine the role of default provisions in a new setting: a one-time, irrevocable choice between two alternative pension plans, a defined benefit (DB) plan and a defined contribution (DC) plan. While DB plans provide retirement benefits that are a set formula based on earnings and service, DC plans provide benefits based on tax-advantaged contributions and subsequent investment performance. Because the plans differ in their accrual patterns and risk characteristics, decisions between the two plans can lead to substantially different retirement income profiles.

We examine a particular firm's transition from a DB plan to a DC plan. While all new employees were offered only the DC plan, the firm offered current workers a one-time opportunity to make a non-reversible choice between the DB plan and the DC plan for all future benefit accruals while employed at the firm. The fact that the choice could not be changed after a particular date represents an important distinction from much of the previous work on defaults in the context of planning for retirement. In particular, decisions regarding asset allocation, the rate of contributions, or even participation in 401(k) plans are choices that affect one's financial circumstances over a horizon as short as a single pay period and can be corrected if mistakes are thought to have been made. By contrast, a non-reversible choice between retirement plans can lead to significantly different amounts of wealth and income at retirement over a worker's career.

In our firm, employees who did not actively choose a retirement plan were defaulted into one plan or the other depending on their age at the time of the transition: individuals age

45 or older were defaulted to remain in the DB plan, while the default for employees under age 45 was to switch to the DC plan. The unique nature of the default, namely that it varied discontinuously by age, allows us to estimate the causal effect of default provisions using a regression discontinuity framework. This methodology differs from past studies of the effect of defaults, which have exploited changes in a default policy regime over time or across localities to identify the effect of the default. While identification in such natural experiments relies on the assumption that individuals exposed to the different regimes, either across time or localities, do not systematically differ in characteristics related to the eventual outcome, in many contexts regression discontinuity relies on weaker identifying assumptions (Imbens and Lemieux 2008).

Figure 1 depicts the percent enrollment in the DC plan as a function of age. By visual inspection, the large discontinuity at age 45 supports the hypothesis that the default had a substantial effect on plan enrollment. Formally, our regression estimates likewise indicate a strong effect: individuals just under the age of 45 are approximately 60 percentage points more likely to choose the DC plan relative to those just over age 45 at the time of plan transition. Given the permanent nature of the decision and the large amount of wealth at stake, the effect is even more dramatic than previous findings of the effects of defaults.

The strong effect we find highlights the importance of setting the default appropriately. We devote the remainder of the paper to deriving a framework for evaluating the potential welfare implications for employees of default provisions and applying it to the firm in our setting. We propose that an optimal default policy maximizes the “aggregate default wealth” or the risk-adjusted value that each employee receives if he or she defaults, aggregated over all employees at the firm. Maximizing the aggregate default wealth is likely to improve employee welfare in many models of household behavior given the strong effects of defaults in this and other contexts.

We examine default rules that are characterized by a cutoff age under which employees are defaulted into the DC plan and where employees older than this cutoff age are defaulted

into the DB plan. This framework generates a representation of age-based default rules for retirement plan transitions that “nudge” employees into plans that raise their expected lifetime utility. We solve for the optimal age cutoff analytically, propose two measures to quantify the gain from introducing an age-based default policy, and provide comparative statics of how the optimal age-based default varies with pension plan, firm, and employee characteristics. We show that the sign of the relationship between the optimal age cutoff and a parameter, such as risk aversion or the contribution rate in the DC plan, depends on how the parameter affects the relative value of the DC plan over the DB plan.

Finally, we illustrate properties of the optimal age cutoff numerically for the firm in our analysis that underwent the transition. We utilize Monte Carlo methods to simulate retirement wealth and variability and employ the constant relative risk aversion family of utility functions. We find that, over a broad range of plausible levels of risk aversion, a heterogeneous default results in a substantially higher aggregate default wealth relative to a uniform policy of defaulting all employees to remain in the DB plan, a common default provision for plan transitions. We also find that an age-based default steers over 99 percent of employees to their higher-value plan, though the employees who appear to be defaulted into their lower-value plan tend to be predominantly women and workers with below-median incomes.

Confirming our analytical results, we find the optimal cutoff age is increasing in the DC contribution rate and decreasing in the generosity of the DB plan. When the level of risk aversion is high, employees generally find the DB plan more attractive which decreases the optimal age cutoff. In addition, the optimal cutoff age depends on the interaction of asset allocation in the DC plan and the level of risk aversion: when the portfolio is invested solely in the most risky asset (i.e. stocks), the optimal cutoff age is higher for low levels of risk aversion, but is lower for high levels of risk aversion. Higher separation rates tend to increase the value of the DC plan more than the DB plan. Finally, the optimal cutoff age is not sensitive to different assumptions regarding wage growth, heterogeneity in risk aversion

by age, or the discount rate, holding all other parameters fixed.

The remainder of the paper is organized as follows. Section 2 provides an overview of the changing pension landscape and describes how our analysis connects the literature on retirement plan choice and default provisions. Section 3 presents the regression discontinuity analysis to identify the causal effect of the default on plan enrollment. Section 4 analytically solves for the optimal age-based default rule and its properties for the general case, while Section 5 empirically illustrates the results from Section 4 using simulated retirement wealth and variability for the employees at the firm in our setting. Section 6 concludes the paper.

2 Background and Related Literature

Over the last 30 years, there has been a pronounced shift from DB plans to DC plans as a result of legislative changes that increased the administrative burden of offering DB plans and changing demand from an increasingly mobile workforce. These changes led employers to replace terminated DB plans with DC plans (Papke 1999; GAO 2008) and made DC plans the most common type of newly implemented employer-provided retirement plan (Kruse 1995). In a recent study, the GAO found that approximately 3.3 million workers who actively participate in a DB plan were affected by a DB plan “freeze,” or the discontinuation of future benefit accruals (GAO 2008). Among these, approximately 83 percent implemented alternative retirement plans, usually a DC plan, implying that a substantial number of employees have experienced a plan migration. A 2003 report by Towers-Perrin indicates that 14 percent of programs that implemented a plan change in recent years allowed current employees to choose between the old and new plan (Towers-Perrin 2003).

The two types of retirement plans differ in the risk faced by participants: participants in DC plans bear the risk of poor investment experience, while DB participants are exposed to the labor market risk of unexpected separations as well as the risk of their employer defaulting on pension obligations. Because individuals are heterogeneous, neither plan dominates the

other, which implies that the transition had a disparate effect on workers.

This paper builds off both the literature on determinants of retirement plan selection and the literature on the role of default provisions on individual's retirement savings behavior. The consensus in the literature on plan selection is that the expected relative value of retirement wealth under either plan depends on an individual's risk preferences, demographic characteristics, and expected job mobility because DB and DC plans differ in how pension wealth accrues (Bodie, Marcus and Merton 1988; Papke 1999; Clark and Pitts 1999; Clark, Ghent and McDermed 2006; Brown and Weisbenner 2007). Studies in this literature evaluate whether the individual determinants of plan selection are consistent with economic theory in terms of which worker characteristics are associated with higher expected retirement wealth for a DC plan relative to a DB plan, or vice versa. These papers use data from institutions, such as public university systems, that offer individuals a choice between enrolling in either a DC or a DB plan (Clark and Pitts 1999; Clark, Ghent and McDermed 2006; Brown and Weisbenner 2007; Manchester *forthcoming*) or those that have changed plan offerings, for example, replacing a DB plan with a DC plan (Papke 2004; Yang 2005).

The second branch of literature – the role of default rules in retirement savings behavior – developed from the fact that DC plans shift the responsibility of enrollment, contribution, and investment decisions to the individual. There is now a large established body of literature that shows that default provisions can have a large impact on savings decisions (for summary, see Beshears, Choi, Laibson and Madrian (2006)). This literature has investigated the role of defaults in decisions regarding participation (e.g., Madrian and Shea 2001), contribution rates (e.g., Choi, Laibson, Madrian and Metrick 2004), asset allocation (e.g., Choi, Laibson and Madrian 2005) and, recently, distributions from DC plans (Mitchell, Mottola, Utkus and Yamaguchi 2009). Typically these studies use changes in 401(k) plan characteristics and default rules at a particular employer as a type of “natural experiment” to investigate the effect of defaults on employee retirement saving, particularly among new hires. This literature has found that despite the fact that a default does not change the menu of options,

switching the default option has substantial effects on retirement savings decisions at each juncture. For example, whether firms have a standard enrollment (i.e. “opt-in”) or an automatic enrollment (i.e. “opt-out”) in 401(k) plans can increase participation rates by more than fifty percentage points (Madrian and Shea 2001). It is important to note, however, that these previous studies estimate changes in an outcome variable that has a limited time horizon because participants can change their enrollment decisions at future dates. On the other hand, plan choice at the time of a transition from DB to DC plans is non-reversible. As pointed out by Beshears et al. (2006), this feature of the decision may mitigate the effect of defaults in the context of plan choice relative to previous results in the literature.

While the role of default rules in savings plans has been well-studied, the literature on the role of default provisions on plan choice is much more limited. In order to determine the causal effect of a default in this context, there must be a choice of plans and variation in the default plan either across employees or over time. Research on firms transitioning from DB plans to DC plans has been in settings where the authors were unable to exploit any variation in the default rule to estimate a causal effect. For example, Papke (2004) and Yang (2005) both examined a plan transition in which there was a one-time transition and a uniform DB default, which precluded the identification of the default’s causal effect on plan choice.

The recent work on defaults in the presence of worker heterogeneity and the possibility of tailoring defaults to individuals provides a key point of connection to the present study. While data limitations stemming from potential legal consequences of treating employees differently have prevented the explicit examination of the effect of incorporating heterogeneity into default provisions (Carrol, Choi, Laibson, Madrian and Mertrick 2009), there are a few examples of heterogeneous defaults in the realm of retirement savings. One such instance is the default distribution from DB pensions, which differs depending on the participant’s marital status: the default distribution for married workers is a joint-and-survivor annuity, while the default for single workers is a single life annuity (Beshears et al. 2006). Another

instance of tailored defaults, which was an outgrowth of the Pension Protection Act (PPA) of 2006, is tying the default asset allocation in DC plans to the participant's age. The PPA included provisions that promoted the use of target-date funds, which set asset allocation based on a specified retirement year range, as qualified default investments. Vanguard (2009) reports that 9 out of 10 plans with automatic enrollment have target-date funds as the default asset allocation, and the default asset allocation in the DC component of the social security systems of Sweden, Chile, and Mexico also depend on the participant's age (Beshears et al. 2006).

In terms of "optimal" defaults, researchers are beginning to consider the importance of worker heterogeneity, particularly when designing and implementing default provisions in the context of plan enrollment. Carrol et al. (2009) find that a third type of enrollment, "active decision," in which individuals must actively choose between enrolling or not is preferable to standard or automatic enrollment even when there is a large degree of heterogeneity in preferred savings rates as long as individuals' short-term discount rates are not too high. When individuals are forced to make an active choice, the effect on enrollment rates in 401(k) plans is comparable to a default policy of automatic enrollment. A recent study by Handel (2009) suggests that removing switching costs by requiring active choice may allow people to make more appropriate health plan choices, but may also lead to worse outcomes in health care markets by increasing adverse selection into health plans.

We contribute to the existing literature in several ways. First, we quantify the effect of default provisions in the context of a non-reversible decision regarding plan choice. Second, we use regression discontinuity to estimate the causal effect of the default, which differs from the methodology used in past studies of defaults. Finally, we provide an analytical framework for solving for the optimal age-based default rule and for assessing the gains from incorporating heterogeneity relative to a uniform default for all employees. We are also able to illustrate these results empirically using data from the firm in our setting.

3 Causal Effect of Default on Plan Enrollment

This section outlines the regression discontinuity methodology used to estimate the causal effect of the default, provides additional details on the employees involved in the transition, and reports the results from estimating the effect of this particular age-based default on plan enrollment.

3.1 Regression Discontinuity Methodology

We implement a standard regression discontinuity framework (see, for example, Imbens and Lemieux (2008) and Lee and Lemieux (2009)) to examine the role of the default rule on plan choice. In particular, we are interested in estimating how having the DC plan as the default relative to having the DB plan affects the probability of switching to the DC plan. In this methodology, the treatment (having the DC plan as the default) is determined by the value of a forcing variable (the employee’s age as of September 1, 2002) being above or below a fixed cutoff value (age 45). Therefore, employees under age 45 received the treatment, while those age 45 or older did not. Because the treatment is a deterministic function of the forcing variable, age, this framework is known as a “sharp” regression discontinuity design. This methodology allows for age to be correlated with plan choice; however, the key assumption is that the relationship between age and the outcome is a smooth, continuous function. This assumption allows any discontinuity at the cutoff value to be interpreted as a causal effect (Imbens and Lemieux 2008). Therefore, the discontinuity in plan enrollment at the cutoff age is the causal effect of the DC default on the probability of enrolling in the DC plan.

Formally, the treatment is the default assignment, given by d_i , and is a deterministic function of the participant’s age as of September 1, 2002, A_i ,

$$d_i = \mathbf{1}\{A_i < c\}, \tag{1}$$

where the variable c denotes the cutoff value and is equal to 45 in this context. A value of

$d_i = 1$ implies that the participant is defaulted into the DC plan and $d_i = 0$ corresponds to the DB plan as the default.

The outcome of interest, enrollment into the DC plan, is given by the variable Y_i . We are interested in the causal impact of d_i on Y_i . This treatment effect, given by τ , is estimated as follows:

$$\tau = \lim_{a \uparrow c} E[Y_i | X_i, A_i = c] - \lim_{a \downarrow c} E[Y_i | X_i, A_i = c] = E[Y_i^{d_i=1} - Y_i^{d_i=0} | X_i, A_i = c], \quad (2)$$

where X_i represents other observable characteristics that are correlated with the outcome variable Y_i . The estimate of τ is thus obtained by estimating the following regression equation:

$$Y_i = \alpha + \beta(A_i - c) + \tau d_i + \gamma(A_i - c)d_i + X_i\pi + \epsilon_i \quad (3)$$

in the interval of $[c - h, c + h]$ for a bandwidth value of h . We examine the sensitivity of our results to different choices of bandwidths and to including higher order powers of $(A_i - c)$.

3.2 Data

To estimate the regression discontinuity estimate of the effect of the default as outlined above, we utilize administrative data from a large non-profit firm that offered 925 existing union employees an opportunity to remain in a DB plan or switch to a DC plan. While the firm employs over 5,000 workers, only existing employees covered by a collective bargaining agreement were eligible for this transition; the existing non-unionized employees were offered a similar one-time opportunity to transition out of the DB plan into the DC plan in 1997.¹ While we do not know the exact motives of the default rule, the firm's decision to discontinue the DB plan to new hires was likely due to increased administrative costs and follows the widespread movement toward DC plans.

¹Data are not available for this earlier transition. However, the default rule for these workers was also an age-based default with a cutoff value of age 45.

These existing employees had approximately six months to make the election. As of September 1, 2002, employees were enrolled in their chosen plan or, if they had failed to make a choice, were defaulted into a plan depending on whether they were older or younger than age 45 on that date. We refer to those who formally chose their plan as “active” participants, and those who were defaulted into a plan as “passive” participants. Table 1 shows the distribution of the type of choice (active vs. passive) and the enrollment decision for the two age groups. Of those employees eligible for the transition, just over half made an active choice. However, the vast majority (70 percent) of employees who made an active choice mimicked the default rule, which is consistent with employees taking the default provisions as advice (Beshears et al. 2006).

After the plan transition, 48 percent of employees were enrolled in the DC plan (Table 2). The dataset contains information regarding the participant’s age, gender, ethnicity/race, hourly wage, tenure, and hours per week. In addition, the employees are divided between two nearby campuses, a primary location where approximately two-thirds of the employees work and a secondary smaller location a few miles away. Summary statistics for these additional variables are provided in Table 2. We restrict the sample to workers under age 65, the normal retirement age at this firm under the DB plan; the average age is approximately 46 years.²

One potential concern is that the sample used for the subsequent empirical evaluation is comprised of only unionized workers who may differ from non-unionized workers. For example, about one-fifth of the sample is female, which is substantially lower than the fraction among the non-unionized employees at this employer. Given that unions have historically favored DB plans over DC plans, the overall participation rates in the DB plan may be higher than in a non-unionized sample; however, the estimation of the causal effect of the default, as outlined in Section 3.1, is not affected by this concern because it is identified from the outcomes of workers on either side of age 45.

²There were 952 existing employees eligible for the transition; the age restriction reduced the sample size to 925.

3.3 Regression Discontinuity Results

In order to validate using regression discontinuity methods to estimate a causal effect, graphical analysis are used to examine: 1) the density of the forcing variable, and 2) the distribution of covariates as a function of the forcing variable. First, we verify that there is no discontinuity in the density of age at the cutoff value of age 45, which would suggest manipulation of the receipt of treatment.³ As shown in Figure 2, there is a relatively constant density of employees surrounding the cutoff value of age 45. Second, we confirm that other observable covariates are a smooth function of the forcing variable and do not experience a discontinuity at the cutoff value as shown in the six panels of Figure 3.

After having verified these conditions for causal inference, we estimate the coefficient τ in Equation (3) by fitting linear probability models to the data assuming a rectangular kernel. Table 3 summarizes regression discontinuity estimates of the effect of the DC default on enrollment into the DC plan assuming a bandwidth of five years. Column (1), the local linear estimate, reports the estimate of τ with no other right-hand side regressors. Column (2) allows the relationship between age and DC enrollment to be linear, with a different slope below and above the discontinuity. Column (3) adds in control variables summarized in Table 2. Columns (4) and (5) provide estimates assuming the relationship between age and DC enrollment follows a cubic function, and differ only in the addition of control variables in Column (5).

These estimates indicate that the default had a strong effect on the enrollment into the DC plan, confirming the initial evidence in Figure 1. An employee just under age 45 was approximately 60 percentage points more likely to enroll in the DC plan than an employee just over age 45. The estimated effect is larger when higher order powers of $(A_i - c)$ are included in the regression. We prefer the estimates in Column (3) because the linear specification fits the data well, as shown in Figure 1, which plots the best fit line and cubic polynomial on either side of the discontinuity. Quantitatively similar results are found

³Given that the forcing variable is based on birth date, manipulation is not an issue in this context.

when estimating the regression using the nonlinear probit model which accounts for the dichotomous nature of the dependent variable. The marginal effects from probit equations are reported in Table 4. The effects are robust to the choice of bandwidth as shown in Appendix A.⁴

Adding demographic control variables does not substantially improve the fit of the model after the default provision is taken into account. Coefficients on gender and race binary variables (not shown) are statistically insignificant, as are coefficient on hours, hourly wage, and whether the employee works in the primary location. Only the coefficient on length of past tenure in the firm is statistically significant, though the estimate is quantitatively small. We also examine whether the causal effect of the default on plan enrollment differs by gender to investigate whether females are more susceptible to the default and we find no difference.⁵

In summary, we find a substantial effect of the default on plan choice: the default increases the probability of enrolling in one plan over the other by 60 percentage points. This effect is, to our knowledge, the first estimate of a default effect in the context of the choice between DB and DC plans. While the magnitude of the estimated effect is comparable to prior work on defaults in the retirement savings literature, the effect estimated here is surprising given the potentially large differences in wealth accrual across the two plans and the non-reversible nature of the decision (Beshears et al. 2006). In addition, the heterogeneity in the default across age provides a unique source of identification for estimating the causal effect of the default, which does not rely on the assumption that employees hired across different default regimes over time are not systematically different from each other as has typically been assumed in existing studies of the effect of defaults.

⁴Cross-validation methods are often used to find the optimal bandwidth in regression discontinuity estimation. However, because the function on either side of the discontinuity is approximately linear, the results are inherently less susceptible to the bandwidth choice (see Figure 1).

⁵These results are not reported but are available upon request.

4 Solving for the Optimal Age-Based Default Rule

The estimates of the default shown in the previous section along with those found in prior literature indicate that defaults can be powerful tools to steer economic behavior. An important question that has been broached in the literature on defaults and in policy circles is how defaults may be constructed to improve welfare. In this section we explore: 1) how to set an optimal default based on observable characteristics; 2) how to determine the value of incorporating heterogeneity into default provisions relative to a uniform policy; and 3) properties of the optimal default as a function of known parameters. Note that the results in this section need not be limited to the particular firm in this paper; in theory, any firm's employee and pension plan characteristics could be utilized to obtain the optimal age-based default rule.

We consider a general specification of expected retirement wealth for DB and DC plans and translate this wealth into certainty equivalents for a general utility function. We then construct an aggregate measure of certainty equivalent that is a function of the cutoff age for a simple age-based default rule. We propose that the optimal cutoff age that maximizes this aggregate measure “nudges” employees toward plans that raise their expected lifetime utility. We evaluate the gain in value from this optimal cutoff age relative to a uniform default policy and the number and characteristics of individual employees who are defaulted into a suboptimal plan under an optimal age-based default policy. Finally, we consider the comparative statics of the optimal cutoff age with respect to different assumptions on the characteristics of the workforce and the employer's retirement plans.

4.1 Characterizing retirement wealth in the DB and DC plan

We begin by characterizing the *future* benefits that accrue in both the DB and the DC plan. While the employee had also accrued years of service prior to the plan migration, the employee receives that benefit stream regardless of the plan they are enrolled in after

the transition. Therefore, the value of past benefits accrued are not included in wealth calculations for either the DB or DC plan and are not relevant for the comparison of future retirement wealth across the two types of plans.

The DB formula is defined by b_j for all years j between 0 and $r - a$, where r is the age of exit from the firm and a is the worker's current age, as a function of annual wages in year j , w_j . The function b_j is completely general in that it can allow for various types of accrual patterns.⁶ The wealth evaluated at retirement age ρ in the DB plan as a function of age a is then:

$$w^{DB}(a) = \int_0^{r-a} b_j(w_j) A_\rho dj, \quad (4)$$

where A_ρ is the actuarial present value of a stream of \$1 annual payments commencing at age ρ and paid until death.

The wealth in the DC plan at retirement age ρ is equal to the contributions made, accumulated with returns from investment experience. Contributions are typically a percentage of annual wages, though they can vary from year to year. We denote employer contributions into the employee's account in year j by c_j and the sequence of returns in all subsequent years by $\delta(k)$ for $k \in [j, \rho - a]$.⁷ The wealth evaluated at retirement age ρ in the DC plan as a function of age a is then:

$$w^{DC}(a) = \int_0^{r-a} c_j w_j e^{\int_j^{\rho-a} \delta(k) dk} dj. \quad (5)$$

We next compute the expected discounted utility of retirement wealth for each plan by explicitly incorporating the uncertainty of different wealth outcomes. We model two sources

⁶For example, a firm that offers 2 percent of average wages for each year of service as the annual retirement benefit would have $b_j(w_j) = 0.02w_j$ for all j . Similarly, a firm that offers its employees 2 percent \times years of service \times the average salary over the last five years at the firm as the annual retirement benefit would be characterized by:

$$b_j(w_j) = \begin{cases} 0 & \text{if } j < r - a - 5 \\ 0.02(r - a) \frac{w_j}{5} & \text{if } j \geq r - a - 5. \end{cases}$$

⁷Note we assume c_j includes only employer contributions, not employee contributions. This distinction is made so both plans' values reflect the benefits provided by the employer only.

of uncertainty: separation risk, which affects the age of exit r , and investment risk, which affects the sequence of returns $\delta(\cdot)$. We assume r and $\delta(\cdot)$ are drawn from a joint distribution $h(r, \delta | a < r \leq \bar{r})$. While separation risk affects the value of the DB and the DC plan, because DB wealth does not depend on $\delta(\cdot)$, investment risk only affects DC retirement wealth.⁸

Assuming a discount rate d , the expected utilities for each plan are given by:

$$EU(w^{DB}(a)) = \int_a^{\bar{r}} \frac{U(w^{DB}(a))}{(1+d)^{\rho-a}} h(r, \delta | a < r \leq \bar{r}) dr \quad (6)$$

$$EU(w^{DC}(a)) = \int_{-\infty}^{\infty} \int_a^{\bar{r}} \frac{U(w^{DC}(a))}{(1+d)^{\rho-a}} h(r, \delta | a < r \leq \bar{r}) dr d\delta. \quad (7)$$

We then define the certainty equivalent for plan $p \in \{DB, DC\}$ as:

$$CE^p(a) = U^{-1}(EU(w^p(a))). \quad (8)$$

The certainty equivalent $CE^p(a)$ for plan p is the amount that makes the individual indifferent between receiving the amount $CE^p(a)$ for certain and the gamble characterized by the uncertain income stream from plan p . Therefore, plan \tilde{p} is preferable to plan \hat{p} if and only if $CE^{\tilde{p}}(a) > CE^{\hat{p}}(a)$.

4.2 Default rules that maximize aggregate default wealth

We posit that the optimal default policy for workers is one that maximizes the “aggregate default wealth.” The aggregate default wealth represents the certainty equivalent that each employee receives if he or she defaults, aggregated over all employees at the firm. Maximizing the aggregate default wealth is likely to improve aggregate welfare because the results in the previous section (and prior research on the effects of defaults) indicate that employees

⁸In our numerical results that follow, we assume independence between separation risk and investment risk. We assume r is drawn from a distribution $f(r | a < r \leq \bar{r}) = \frac{f(r)}{1-F(a)}$ where F is the cumulative distribution function of f and that the sequence $\delta(\cdot)$ is determined by a draw from a distribution $g(\delta)$. However, this assumption does not affect our analytical results, which are all in terms of the certainty equivalent.

are likely to enroll in their default plan, either due to inertia, transaction costs, complexity, procrastination, or the perception of endorsement (Beshears et al. 2006). Further justification for this objective function comes from the fact that the cost of a worker enrolling in the plan that is not the default plan is low. Therefore, because workers who determine which plan is best can easily enroll in their chosen plan even if it is not their default, focusing on maximizing the welfare of those who default seems appropriate. While in this section, we focus on policies that are likely to maximize employee welfare, we recognize that firms may have alternative objective functions that take into account the relative costs of the two plans. We revisit this subject briefly in Section 4.3.

We consider age-based default policies that take the following form for a given cutoff age \tilde{a} : 1) individuals younger than \tilde{a} are defaulted into the DC plan, and 2) individuals older than \tilde{a} are defaulted into the DB plan. We solve for the optimal age-based default policy that maximizes the aggregate default wealth under the constraints outlined above.

Define the minimum age among employees at the firm as \underline{a} and the maximum age as \bar{a} . Formally, the optimal age-based default policy that satisfies conditions (1) and (2) above is defined by a^* , where a^* maximizes:

$$\begin{aligned} L(a^*) &= \int_{\underline{a}}^{a^*} CE^{DC}(a) da + \int_{a^*}^{\bar{a}} CE^{DB}(a) da \\ &= \int_{\underline{a}}^{\bar{a}} CE^{DB}(a) da + \int_{\underline{a}}^{a^*} (CE^{DC}(a) - CE^{DB}(a)) da. \end{aligned} \quad (9)$$

The first term of Equation (9) does not depend on a^* . Therefore we can solve for a^* as:

$$a^* = \arg \max_{\tilde{a}} \int_{\underline{a}}^{\tilde{a}} (CE^{DC}(a) - CE^{DB}(a)) da. \quad (10)$$

Solving the problem in Equation (10) yields the first-order condition for an interior solution:

$$H(a^*; \gamma) \equiv (CE^{DC}(a^*) - CE^{DB}(a^*)) = 0, \quad (11)$$

where γ is the vector of parameters $[\gamma_l]$ that define the solution a^* .

By the implicit function theorem, we can derive the comparative statics of a^* as follows:

$$\frac{\partial a^*}{\partial \gamma_l} = -\frac{H_{\gamma_l}(a^*; \gamma)}{H_{a^*}(a^*; \gamma)} = -\frac{\frac{\partial CE(a)^{DC}}{\partial \gamma_l} - \frac{\partial CE(a)^{DB}}{\partial \gamma_l}}{\frac{\partial CE(a)^{DC}}{\partial a} - \frac{\partial CE(a)^{DB}}{\partial a}} \Bigg|_{a=a^*}. \quad (12)$$

Because a^* maximizes $L(a^*)$, the denominator in Equation (12) is negative. Therefore,

$$\text{sign} \left(\frac{\partial a^*}{\partial \gamma_l} \right) = \text{sign} \left(\frac{\partial CE(a)^{DC}}{\partial \gamma_l} - \frac{\partial CE(a)^{DB}}{\partial \gamma_l} \right) \Bigg|_{a=a^*}. \quad (13)$$

Intuitively, Equation (13) shows that the optimal age-based default rule defined by a^* is increasing in parameters that improve the relative value of the DC plan over the DB plan, which would then default additional employees into the DC plan. For example, an increase in the generosity of the employer contribution to the DC plan would lead to an increase in the optimal cutoff age. Similarly, a more generous DB benefit plan formula would lead to a decrease in the optimal cutoff age. More details on comparative statics can be found in Appendix B.

The ability to set the default based on an observable characteristic rather than having the same default for all employees at the firm is at worst welfare-neutral since the choice of setting the default to be equal for all employees (i.e., a^* equal to \bar{a} or \underline{a}) is a choice available under the optimization procedure outlined above. But how much do employees gain from a heterogeneous default that differs by age?

We quantify the gain in aggregate default wealth from an age-based default policy by the measure G , which represents the percent increase in the aggregate default wealth for the age-based default rule relative to a commonly used default rule, one that defaults everyone to remain in the DB plan. The measure G is defined as:

$$G \equiv \frac{L(a^*) - L(\underline{a})}{L(\underline{a})}, \quad (14)$$

where the function L defined in Equation (9) denotes the aggregate default wealth. $L(a^*)$ denotes the aggregate default wealth under the optimal age-based default rule and $L(\underline{a})$ denotes the aggregate default wealth for a uniform DB default.

While the measure G may show that the firm's workers as a whole are better off under a heterogeneous default policy, it does not give any sense of whether an age-based default policy accurately categorizes workers into the plan that maximizes each worker's certainty equivalent. In particular, because each worker's certainty equivalent depends not only on age, but also on gender and wages, it is possible that conditioning the default only on age defaults some workers into the plan with a lower certainty equivalent. We therefore also determine the number of employees who fall in such a category, N_l , and the average loss in certainty equivalent for these individuals, l , as follows:

$$N_l \equiv \int_{\underline{a}}^{a^*} \mathbf{1}_{[CE^{DC}(a) < CE^{DB}(a)]} da + \int_{a^*}^{\bar{a}} \mathbf{1}_{[CE^{DC}(a) > CE^{DB}(a)]} da \quad (15)$$

$$l \equiv \left(\int_{\underline{a}}^{a^*} (CE^{DB}(a) - CE^{DC}(a)) \mathbf{1}_{[CE^{DC}(a) < CE^{DB}(a)]} da + \int_{a^*}^{\bar{a}} (CE^{DC}(a) - CE^{DB}(a)) \mathbf{1}_{[CE^{DC}(a) > CE^{DB}(a)]} da \right) / N. \quad (16)$$

In our numerical simulation, we provide some descriptives on the individuals who are defaulted into a suboptimal plan under the optimal age-based default rule.

We explicitly only model separation risk and investment risk above and ignore the risk that the employer goes bankrupt and can no longer honor DB pension obligations and the risk of dying prior to retirement. Nonetheless, it can easily be shown that the results above are general to any types of risk. It is important to note, however, that our results implicitly assume that utility is additively separable in other sources of retirement wealth (for example, from past or future employer plans or personal savings accounts) and ignores the potential value of other benefits from DB or DC plans, such as the ability to borrow against DC plan

balances. We also ignore any potential effects on other components of total compensation, for instance, the possibility that more workers enrolling in the plan that gives them higher benefits increases the cost to the employer, thereby causing the employer to reduce wages or other fringe benefits. Despite these limitations, our framework provides valuable insights into how heterogeneity may be introduced into default provisions.

4.3 Default rules that account for firm costs

The previous section assumed that the firm's objective was to maximize employee welfare. While we believe that this objective is relevant for policy considerations, we acknowledge that firms may also take into account costs when making decisions related to employee benefits.

As an alternative, suppose firms operate in a two-stage manner where they choose how much to spend on deferred compensation in the first stage and then choose a policy to maximize employee welfare subject to this constraint in the second stage. Our maximization problem in Equation (10) can be easily modified to accommodate this optimization strategy as follows:

$$a^{**} = \arg \max_{\tilde{a}} \int_{\underline{a}}^{\tilde{a}} (CE^{DC}(a) - CE^{DB}(a)) da \quad (17)$$

subject to the firm budget constraint:

$$\int_{\underline{a}}^{a^{**}} FC^{DC}(a) da + \int_{a^{**}}^{\bar{a}} FC^{DB}(a) da \leq B, \quad (18)$$

where $FC^p(a)$ denotes net firm costs for plan p for a worker age a and B is the firm's budget for deferred compensation. Net firm costs could include the present value of future retirement benefits offset by any benefit fueled by the retirement plan, e.g. differences in turnover costs across the two plans. This modified problem is equivalent to the original problem if the budget constraint is not binding, i.e., $B \geq \int_{\underline{a}}^{a^*} FC^{DC}(a) da + \int_{a^*}^{\bar{a}} FC^{DB}(a) da$, where a^* denotes the optimal cutoff found in Section 4.2. If the budget constraint is binding, the optimal cutoff would be chosen from the feasible set of cutoff policies.

A second alternative is to consider a social planner's problem which maximizes total surplus. A social planner would choose to find the optimal cutoff policy that minimizes overall benefits less costs:

$$a^{***} = \arg \max_{\tilde{a}} \int_{\underline{a}}^{\tilde{a}} [(CE^{DC}(a) - FC^{DC}(a)) - (CE^{DB}(a) - FC^{DB}(a))] da. \quad (19)$$

The first order condition of the social planner's problem equates marginal benefits to marginal costs:

$$CE^{DC}(a^{***}) - FC^{DC}(a^{***}) = CE^{DB}(a^{***}) - FC^{DB}(a^{***}). \quad (20)$$

In summary, an optimal age-based default rule may be obtained if the firm's objective differed from one which only considered employee welfare.

In the following section, we illustrate the solution to the problem which maximizes employee welfare as described in Section 4.2. We do this for two main reasons. First, we believe the firm in our setting had the objective of maximizing the welfare of existing employees at the time of the plan transition. The non-profit nature of the firm gives rise to a lower emphasis on a policy that minimizes the cost to the firm, and it is plausible that the firm is sympathetic to the preferences of existing workers who were hired under the DB regime. The strong effect of the default that we find in Section 3 is consistent with this idea because the employees may have been more reluctant to follow the default if it perceived the firm's motivation as purely cost minimization.

Second, solving for the optimal age cutoff by solving Equation (17) or (19) requires a method for calculating firm costs. Firm costs would include not only the present value of future contributions into the DC plan and benefits paid from the DB plan, but would also depend on the uncertainty in funding DB benefits, administrative costs to either plan, and the firm's cost of raising capital. While we do not have the ability of incorporating these components, if they were known, the optimization procedure could be modified accordingly.

5 Numerical Simulations of the Optimal Age-Based Default Rule

In this section we empirically evaluate the optimal age-based default using data from the large, non-profit firm that was the focus of the analysis in Section 3. This allows us to evaluate how the optimal cutoff age compares to the firm’s chosen cutoff age of 45. We quantify the gain in the aggregate default wealth from incorporating a simple age-based default relative to a uniform default policy, and provide descriptive summaries of the employees who are defaulted into their lower-value plan. Finally, we illustrate the comparative statics, or how the optimal cutoff changes with values of the parameters.

5.1 Simulation methods and assumptions

We discretize the optimization procedure in the previous section and consider age-based default policies where a^* is an integer. The firm’s DB formula provides workers a stream of payments equal to a constant percentage of the employee’s average wage for each year of service at the firm, so $b_j(w_j) = bw_j$ for all j . Therefore, the wealth evaluated at retirement age ρ for the DB plan is:

$$w^{DB}(a) = \sum_{j=1}^{r-a} bw_j a_\rho. \quad (21)$$

The wealth from the DC plan evaluated at retirement age ρ assuming a constant contribution rate $c_j = c$ for all j is:

$$w^{DC}(a) = \sum_{j=1}^{r-a} cw_j \prod_{k=j}^{\rho-a-1} (1 + r_{sk}), \quad (22)$$

where s represents a draw from a distribution of investment return sequences.

We model uncertainty in r , the exit age from the firm, and uncertainty in the sequence of investment returns. Separation probabilities are assumed to follow a constant hazard

rate by age, and summarized by probabilities $p_r(a)$ where $\sum_{r=a+1}^{\rho} p_r(a)$ for each a . We also simulate S draws of investment return sequences using a Monte Carlo method outlined and used by Shoven and Sialm (1998, 2003) that draws series of asset returns for three different asset classes: stocks, bonds, and money market. The distribution of returns is assumed to be lognormal and draws are assumed to be serially-independent.⁹ Our numerical results assume that separation risk and investment risk are independent.

The expected discounted utilities for the DB and the DC plans are given by:

$$EU(w^{DB}(a)) = \sum_{r=a+1}^{\rho} p_r(a) \frac{U(w^{DB}(a))}{(1+d)^{\rho-a}} \quad (23)$$

$$EU(w^{DC}(a)) = \sum_{s=1}^S \frac{1}{S} \sum_{r=a+1}^{\rho} p_r(a) \frac{U(w^{DC}(a))}{(1+d)^{\rho-a}}. \quad (24)$$

The assumptions used to simulate DB and DC retirement wealth are based on the known characteristics of the firm in our setting and are reported in Table 5. The DB multiplier b is equal to 2 percent at this firm. The employer's contribution to the DC plan is comprised of 5 percent of the employees' annual salary plus additional matching contributions, up to a total contribution of 10 percent. We model our results assuming the median match percentage of 3.5 percent of salary for a total employer contribution rate c equal to 8.5 percent. In evaluating both the DB and DC retirement wealth, we assume real wages grow at a constant 2 percent per year, and a 2.5 percent rate of inflation. We assume a constant real discount rate d of 1 percent, and a constant separation hazard of 5 percent, taken from the data. The normal retirement age at the firm is 65, so we assume $\rho = 65$. Annuity values are taken from Social Security Administration mortality tables for males and females born in 1960 and assume a 2.9 percent interest rate, implying an annuity value of 14.48 for women and 13.15 for men. The mean, standard deviations, and covariances for different asset classes used in the simulation are calibrated based on historical real returns reported by Ibbotson Associates (Ibbotson 2008) and summarized in the bottom panel of Table 5. We assume

⁹For more details regarding the simulation methods, please see Appendix C.

asset allocation follows the pattern of Fidelity Investments target-date retirement funds, the default fund allocation for DC participants at this firm, as shown in Figure 4.¹⁰ When illustrating the comparative statics of the optimal age-based default policy, we vary each of these parameters from their baseline values independently.

We use the standard constant relative risk aversion (CRRA) utility function to evaluate the expected discounted utility for each plan and to then translate it into a certainty equivalent for each individual as in Equation (8). The functional form of the utility function is given by:

$$U(w) = \frac{w^{1-\alpha}}{1-\alpha}, \quad (25)$$

where α is the measure of relative risk aversion. Estimates of α in the literature vary from 1 to 10, or higher depending on the context (e.g., Mehra and Prescott 1985; Kocherlakota 1996; Chetty 2006). Recent work by Goldstein, Johnson and Sharpe (2008) estimate $\alpha = 6.1$ from an experimental setting in which participants indicated their preferences for a distribution of future retirement income given a cost constraint. In our empirical analysis we vary α from 0 to 10 to evaluate how a^* changes with the level of risk aversion among employees at the firm.

5.2 Main simulation results

The panels in Figure 5 show the certainty equivalent by age for each pension plan across different levels of risk aversion. The absolute levels of the certainty equivalent wealth are decreasing in the relative risk aversion, ranging from an average of \$112,106 (DC) and \$98,390 (DB) for $\alpha = 0$ to \$7,261 (DC) and \$18,015 (DB) for $\alpha = 10$. For lower levels of risk aversion, the certainty equivalent wealth is higher in the DC plan for younger workers as shown in panels (a) and (b) in Figure 5. For the highest level of risk aversion shown (panel (d)), the certainty equivalent from the DB plan exceeds the DC plan for all ages. The panels in

¹⁰We use the allocation for ages between 20 and 65 in 5-year intervals assumed by current funds to fit a fractional multinomial logit model with fourth order age terms in order to estimate an implied asset allocation across the three classes for ages between the 5-year intervals.

Figure 6 plot the difference between the certainty equivalent wealth across the two plans, or $CE_i^{DC} - CE_i^{DB}$. The graphs show at approximately which age the certainty equivalents of the DB plan are equated to the certainty equivalents of the DC plans, which is where the plots cross the x -axis.

As outlined in Section 4, we define the optimal cutoff age a^* for the simple, age-based default as the one which maximizes the aggregate default wealth (Equation (9)). We use our baseline parameter assumptions to solve for a^* for different levels of risk aversion. The objective function in Equation (9) is plotted against age for different values of α in the panels of Figure 6. The age at which the objective function is maximized is a^* , and is reported in Table 6 for different values of α . The optimal cutoff value ranges from 44 for $\alpha = 2$ to 20 for $\alpha = 10$, and is non-monotonic in nature due to the opposing sources of risk: the presence of investment risk implies that more risk-averse individuals would prefer the DC plan, but separation risk affects the value of the DB plan more than the DC plan and thus pushes the optimal cutoff age in the opposite direction. For levels of risk aversion less than 4, the optimal age cutoff is between 42 and 47, largely in line with the firm's chosen age cutoff of 45. Therefore, our simulation results indicate that the firm's cutoff age is close to optimal for low levels of risk aversion under the assumption that their objective was to maximize employee welfare. The lowest value for α such that a uniform DB default is optimal is approximately equal to 8. Therefore, we find that for a sizable range of risk aversion, an age-based default rule produces a higher aggregate default wealth than a uniform DB default.

We quantify the gain in aggregate default wealth using the measure G as outlined in Equation (14) and report the results in Tables 6. The value G in the second row of Table 6 compares the aggregate default wealth under the optimal age-based default to the aggregate default wealth under a uniform DB default. We find that implementing a simple age-based default leads to a substantial increase in the aggregate default wealth under the optimal policy. Because employees are predisposed to choosing the default plan, as shown in Section 3, it is likely that a heterogeneous default would lead to a substantial increase in welfare.

The measure G masks the fact that certain employees may get defaulted into the plan with the lower certainty equivalent under the outlined optimization procedure. Table 6 summarizes the number of people for whom the certainty equivalent is lower under the optimal default policy than the alternative, as outlined in Equations (15) and (16). Less than 1 percent of employees are defaulted into the lower-value plan, amounting to an average loss of approximately \$4,000 in certainty equivalent (or less than 5 percent) for lower levels of risk aversion. All of the workers who lose from the optimal default policy are women, and the majority are workers with below-median income. This finding is likely due to the low weight that women and lower income workers receive in the aggregate default wealth due to their smaller numbers (in the case of women) or lower certainty equivalents (for lower-income workers). Weighting functions that counteract this underrepresentation could address this issue.

5.3 Sensitivity analysis and robustness to alternative assumptions

In Table 8 we illustrate how the optimal age cutoff changes with different assumptions regarding pension plan characteristics, asset allocation, separation hazards, investment risk, wage growth, the discount rate, and inflation. Our results confirm the general predictions given in Section 4: changes in the parameter values that increase the value of the DC plan relative to the DB plan increase the optimal cutoff age.

We find that increasing the generosity of the DC contribution rate or reducing the generosity of the DB plan increases the optimal cutoff age across all values of α . The optimal cutoff age depends on the asset allocation chosen in the DC plan such that when the portfolio is invested solely in stocks, the riskiest asset class, the optimal cutoff age is higher for lower levels of risk aversion and is lower for higher levels of risk aversion. Analogously, investing only in bonds reduces the optimal cutoff age; however, a^* does not vary as dramatically across different levels of risk aversion. Under the assumption that DC participants are invested solely in low-yielding risk-free assets, the DB plan performs better even under high

levels of risk aversion.

Eliminating separation risk entirely illustrates that isolating investment risk produces a monotonically decreasing optimal cutoff age. Similarly, eliminating investment risk while maintaining the base assumption for separation risk yields a monotonically increasing optimal cutoff age. As separation risk increases, the certainty equivalent in the DB plan is reduced more than the certainty equivalent in the DC plan, and, therefore, the optimal cutoff age increases, defaulting additional employees into the DC plan. If the standard deviations of the three asset classes are doubled, the optimal age cutoff increases for low levels of risk aversion and decreases for high levels of risk aversion relative to the baseline.¹¹

The optimal cutoff age is weakly decreasing in real wage growth, indicating that higher wage growth increases the relative value of DB plan benefits over DC plan benefits, though the effect is small. The real discount rate does not affect the relative values of the DB and DC certainty equivalents; therefore, the optimal cutoff age does not change with respect to the real discount rate.

The inflation rate assumption affects the optimal cutoff age because inflation differentially affects the value of the DB relative to the DC plan. DB plan participants' benefits are based on a formula that includes their nominal wages, from what may be many years prior to retirement. By contrast, early contributions in DC accounts are invested in markets that implicitly adjust for inflation. High inflation therefore reduces the value of DB benefits relative to DC benefits and increases the optimal cutoff age.¹²

We examine how the optimal age cutoff changes if there is heterogeneity in risk aversion

¹¹Note these results assume that investment risk and separation risk are independent. If we assume that they are correlated, perhaps because poor economic conditions reduce investment returns and increase involuntary separations, these two sources of risk would tend to counteract each other, as high separation risk reduces the relative value of the DB plan and low investment returns reduce the relative value of the DC plan.

¹²Note that the low sensitivity to wage growth and the higher sensitivity to the inflation rate is likely due to the fact that our firm's DB benefit formula uses wages throughout a worker's career rather than common alternatives, such as a formula that bases benefits on final average salary. This type of formula would be more likely to protect DB participants from inflation and would reward higher wage growth more than a DC plan. On the other hand, the relative effect of separation risk on the DB plan is likely to be greater with a DB formula that uses a final average salary.

across employees. The literature shows evidence of a small positive correlation between risk aversion and age (Barsky, Juster, Kimball and Shapiro (1997), Kimball, Sahn and Shapiro (2009)) but there is not evidence that risk aversion is strongly related to gender (see ? for review). In the last column of Table 8, we vary the risk aversion parameter such that the average risk aversion for employees is equal to 2, but risk aversion increases with age linearly between 1.1 and 2.6. The optimal cutoff ages under this assumption are very similar to those in Column 2 where the employees all exhibit a risk aversion parameter of 2. We conclude that our results are insensitive to a positive correlation between age and risk aversion.¹³

In our analysis, we do not explicitly value the differences in death benefits and vesting requirements between the two plans. In our firm, the DB plan provides a survivor benefit to a named beneficiary equal to 50 percent of the accrued retirement benefit, while the DC plan provides the survivor the entire worker’s DC account. This makes the DB plan less attractive for all workers, with a greater differential for younger workers who have a higher likelihood of dying at some point before retirement relative to older workers who are closer to retirement. Similarly, because DB benefits vest after 5 years of service while employer contributions to DC accounts vest after only 1 year, incorporating vesting requirements into our analysis would have an analogous effect, raising the relative value of the DC plan for all workers and particularly for younger workers who are more likely to have less tenure at the firm. The combination of these effects would slightly increase the optimal cutoff age, defaulting more workers into the DC plan.

6 Conclusion

The idea that individual choices can depend on institutional arrangements, such as the default provisions that dictate what happens when an individual fails to make an active decision, is now widespread in economics. This idea has given rise to a notion of “libertarian

¹³This statement holds for different average levels of risk aversion and different amounts of dispersion by age.

paternalism,” an approach that encourages private and public bodies to steer individuals to choices that are likely to make them better off while still preserving freedom of choice (Thaler and Sunstein 2003). The current study provides strong evidence that a firm can significantly influence the choice of a pension plan by its decision of which plan is the default: our estimates indicate that employees who were subject to a DC plan default were 60 percentage points more likely to enroll in a DC plan relative to a DB plan. The default was an overwhelming determinant of plan enrollment, alone explaining a substantial amount of the variability in plan choice. This evaluation differs from past studies of defaults in that we examine a one-time, non-reversible decision with potentially substantial retirement wealth at stake and we are able to use a regression discontinuity framework to estimate the causal effect of the default rule.

Given that we find a strong effect of the default on plan enrollment, we develop a framework for finding the optimal age-based default rule for pension plan choice. We solve for the age cutoff that maximizes the aggregate default wealth, or the aggregate certainty equivalent that each employee receives upon failing to make an active decision. We posit that maximizing this aggregate default wealth is likely to improve aggregate employee welfare given the strong effects of defaults in our context as well as in other settings.

We numerically simulate the optimal cutoff age and the gain in the aggregate default wealth from introducing an age-based default policy. We find that over a broad range of plausible levels of risk aversion, a heterogeneous default produces a substantially higher aggregate default wealth relative to a uniform policy of defaulting all employees into a DB plan, a common default policy used by firms transitioning from a DB plan to a DC plan. Our results also suggest that age is an important determinant of the relative value of a DC plan over a DB plan, and that incorporating observable characteristics in addition to age into the determination of the default would not significantly alter employee outcomes. We also demonstrate how the optimal age-based default rule varies with different levels of risk aversion, asset allocations strategies, and plan characteristics.

It is important to note that the optimal age cutoff is sensitive to the level of risk aversion among employees at the firm, and that incorporating risk aversion into a heterogeneous default policy would likely produce outcomes that are welfare-enhancing. However, such an approach may be difficult to implement. If employers systematically under- or over-estimated the level of risk aversion among their employees when constructing an optimal age-based default policy, or if there were substantial heterogeneity in risk aversion across employees, it is possible that more employees could be defaulted into a suboptimal plan. However, there have been advances in survey methods to assess risk aversion that could be utilized by firms to learn about their employees' levels of risk aversion, and we show that the optimal age cutoff is not sensitive to incorporating heterogeneity in risk aversion across observable characteristics, such as age.

Overall, our results suggest that substantial welfare gains are possible by varying defaults by observable characteristics. In particular, when a set of observable characteristics strongly predicts the relative value of one choice over others, conditioning a default on these characteristics can “nudge” individual decision-makers into choices that are likely to raise their expected lifetime utility.

References

- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro**, “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *Quarterly Journal of Economics*, 1997, *112* (2), 537–579.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian**, “The Importance of Default Options for Retirement Savings Outcomes: Evidence from the United States,” *NBER Working Paper No. 12009*, 2006.
- Bodie, Zvi, Alan J. Marcus, and Robert C. Merton**, “Defined Benefit versus Defined Contribution Pension Plans. What Are the Real Trade-offs?” in Zvi Bodie, John Shoven, and David Wise, eds., *Pensions in the U.S. Economy*, Chicago, IL: University of Chicago Press, 1988.
- Brown, Jeffrey R. and Scott J. Weisbenner**, “Who Chooses Defined Contribution Plans?,” *NBER Working Paper No. 12842*, 2007.
- Carrol, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Mertrick**, “Optimal Defaults and Active Decisions,” *Quarterly Journal of Economics*, 2009, *124* (4), 1639–74.
- Chetty, Raj**, “A New Method of Estimating Risk Aversion,” *American Economic Review*, 2006, *96* (5), 1821–1834.
- Choi, James J., David Laibson, and Brigitte C. Madrian**, “Are Empowerment and Education Enough? Underdiversification in 401(k) Plans,” *Brookings Papers on Economic Activity*, 2005, (2), 151–198.

- , — , — , and **Andrew Metrick**, “For Better or For Worse: Default Effect and 401(k) Savings Behavior,” in David A. Wise, ed., *Perspectives in the Economics of Aging*, Chicago, IL: University of Chicago Press, 2004.
- Clark, Robert L. and M. Melinda Pitts**, “Faculty Choice of Pension Plan: Defined Benefit versus Defined Contribution,” *Industrial Relations*, 1999, 38 (1), 18–45.
- , **Linda S. Ghent, and Ann A. McDermed**, “Pension Plan Choice among University Faculty,” *Southern Economic Journal*, 2006, 72 (3), 560–577.
- DellaVigna, Stefano**, “Psychology and Economics: Evidence from the Field,” *Journal of Economic Literature*, 2009, 47 (2), 315–372.
- GAO**, “Defined Benefit Pensions: Plan Freeze Affect Millions of Participants and May Post Retirement Income Challenges,” Report Number: GAO-08-817, General Accounting Office, Washington DC 2008.
- Goldstein, Daniel, Eric J. Johnson, and William F. Sharpe**, “Choosing Outcomes versus Choosing Products: Consumer-Focused Retirement Investment Advice,” *Journal of Consumer Research*, October 2008, 35, 440–456.
- Handel, Benjamin**, “Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts,” mimeo., Northwestern University 2009.
- Ibbotson**, *Stocks, Bonds, Bills, and Inflation 2008 Yearbook*, Chicago, IL: Ibbotson Associates, 2008.
- Imbens, Guido and Thomas Lemieux**, “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, 2008, 142 (2), 615–635.
- Johnson, Eric J. and Daniel Goldstein**, “Do Defaults Save Lives?,” *Science*, 2003, 302, 1338–1339.

- Kimball, Miles S., Claudia R. Sahn, and Matthew D. Shapiro**, “Risk Preferences in the PSID: Individual Imputations and Family Covariation,” *American Economic Review*, May 2009, *99* (2), 363–368.
- Kocherlakota, Narayana R.**, “The Equity Premium: It’s Still a Puzzle,” *Journal of Economic Literature*, 1996, *34* (1), 42–71.
- Kruse, Douglas L.**, “Pension Substitution in the 1980s: Why the Shift toward Defined Contribution?,” *Industrial Relations*, 1995, *34* (2), 218–241.
- Lee, David S. and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *NBER Working Paper No. 14723*, 2009.
- Madrian, Brigitte C. and Dennis F. Shea**, “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior,” *Quarterly Journal of Economics*, 2001, *116* (4), 1149–1525.
- Manchester, Colleen F.**, “The Effect of Pension Plan Type on Retirement Age: Distinguishing Plan Incentives from Preferences,” *Southern Economic Journal*, *forthcoming*.
- Mehra, Rajnish and Edward C. Prescott**, “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 1985, *15*, 145–161.
- Mitchell, Olivia S., Gary R. Mottola, Stephen P. Utkus, and Takeshi Yamaguchi**, “Default, Framing, and Spillover Effects: The Case of Lifecycle Funds in 401(K) Plans,” NBER Working Paper 15108, National Bureau of Economic Research 2009.
- Papke, Leslie**, “Are 401(k) Plans Replacing Other Employer-Provided Pensions? Evidence from Panel Data,” *Journal of Human Resources*, 1999, *34* (2), 346–369.
- , “Pension Plan Choice in the Public Sector: The Case of Michigan State Employees,” *National Tax Journal*, 2004, *57* (2), 329–339.

Shoven, John B. and Clemens Sialm, “Long Run Asset Allocation for Retirement Savings,” *Journal of Private Portfolio Management*, 1998, pp. 13–26.

— and —, “Asset Location in Tax-Deferred and Conventional Savings Accounts,” *Journal of Public Economics*, 2003, 88, 23–38.

Thaler, Richard H. and Cass R. Sunstein, “Libertarian Paternalism,” *American Economic Review*, May 2003, 93 (2), 175–179.

Towers-Perrin, “Back to the Future: Redefining Retirement in the 21st Century,” 2003.

Vanguard, “How American Saves 2009,” 2009.

Yang, Tongxuan, “Understanding the Defined Benefit versus Defined Contribution Choice,” Pension Research Council Working Paper Series 2005-04, The Wharton School, University of Pennsylvania 2005.

Table 1: Distribution of Age Group by Choice (*Active vs. Passive Choice*)

Age Group (<i>default</i>)	Passive Choice	Active Choice	
		Consistent with Default	Inconsistent with Default
Under 45 (<i>DC Plan</i>)	194	162	41
45 or Over (<i>DB Plan</i>)	308	134	86
Total	502	296	127

Notes: Sample restricted to employees less than 65 years of age. N=925.

Table 2: Summary Statistics: Employee Characteristics

Variable	Mean	Std. Dev.	Min	Max
Enrolled in DC Plan	0.478	0.500	0	1
Made Passive Choice	0.543	0.498	0	1
Age	46.07	9.72	21.88	64.96
Female	0.188	0.391	0	1
White	0.425	0.495	0	1
Black	0.114	0.317	0	1
Hispanic	0.303	0.460	0	1
Other Race/Ethnicity	0.159	0.366	0	1
Weekly Hours	39.35	3.27	20.00	55.00
Hourly Wage	23.98	6.63	10.24	36.81
Tenure (years)	12.08	9.12	0.75	43.41
Primary Work Location	0.699	0.459	0	1

Notes: Sample restricted to employees less than 65 years of age. N=925.

Table 3: Regression Discontinuity Estimates of Default Effect on Plan Choice (*Linear Probability Model*)

	(1)	(2)	(3)	(4)	(5)
Under 45	0.605*** (0.042)	0.576*** (0.080)	0.578*** (0.081)	0.777*** (0.141)	0.801*** (0.143)
(Age - 45)		-0.017 (0.020)	-0.011 (0.020)	0.369** (0.181)	0.437** (0.179)
(Age - 45) ²				-0.169** (0.085)	-0.198** (0.084)
(Age - 45) ³				0.020* (0.011)	0.024** (0.011)
(Age - 45)×Under 45		0.025 (0.029)	0.023 (0.029)	-0.4 (0.266)	-0.462* (0.261)
(Age - 45) ² ×Under 45				0.129 (0.130)	0.167 (0.130)
(Age - 45) ³ ×Under 45				-0.028 (0.018)	-0.030* (0.018)
Hours			-0.004 (0.005)		-0.005 (0.006)
Hourly Wage			0.005 (0.005)		0.006 (0.005)
Tenure			-0.006** (0.003)		-0.007** (0.003)
Primary Work Location			0.068 (0.065)		0.075 (0.065)
Constant	0.227*** (0.030)	0.272*** (0.061)	0.356 (0.268)	0.077 (0.101)	0.113 (0.302)
R^2	0.359	0.357	0.365	0.358	0.368
N	353	353	353	353	353

Notes: Dependent variable is enrollment in the DC plan. Regression discontinuity estimate is coefficient on “Under 45,” which estimates the change in DC plan enrollment at the age cutoff. Robust standard errors in parentheses. Bandwidth of 5 years. Columns (3) and (5) also include gender and race binary variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Regression Discontinuity Estimates of Change in DC Plan Enrollment at Default Age (*Probit Model*)

	(1)	(2)	(3)	(4)	(5)
Under 45 (d)	0.605*** (0.042)	0.581*** (0.082)	0.601*** (0.083)	0.758*** (0.116)	0.784*** (0.110)
(Age - 45)		-0.023 (0.027)	-0.017 (0.027)	0.523* (0.271)	0.644** (0.273)
(Age - 45) ²				-0.239* (0.124)	-0.291** (0.125)
(Age - 45) ³				0.029* (0.016)	0.035** (0.016)
(Age - 45) × Under 45		0.035 (0.043)	0.038 (0.045)	-0.578 (0.421)	-0.727* (0.426)
(Age - 45) ² × Under 45				0.172 (0.205)	0.22 (0.210)
(Age - 45) ³ × Under 45				-0.041 (0.028)	-0.047 (0.029)
Hours			-0.011 (0.011)		-0.011 (0.012)
Hourly Wage			0.007 (0.007)		0.009 (0.007)
Tenure			-0.010** (0.005)		-0.010** (0.005)
Primary Work Location (d)			0.102 (0.090)		0.113 (0.089)
R^2	0.28	0.281	0.308	0.291	0.32
N	353	353	353	353	353

Notes: Marginal effects reported; (d) denotes marginal effects computed for discrete change of binary variable from 0 to 1. Dependent variable is enrollment in the DC plan. Regression discontinuity estimate is marginal effect of “Under 45,” which estimates the change in DC plan enrollment at the age cutoff. Robust standard errors in parentheses. Bandwidth of 5 years. Columns (3) and (5) also include gender and race binary variables. * p<0.10, ** p<0.05, *** p<0.01.

Table 5: Baseline Assumptions for Optimal Age-Based Default

	Assumption	
<i>Plan Characteristics:</i>		
DB Multiplier (b)	2.0%	
DC Contribution Rate (c)	8.5%	
<i>Other Parameters:</i>		
Real Wage Growth Rate (g)	2.0%	
Real Discount Rate (d)	1.0%	
Separation Hazard ($\Rightarrow p_r^a$)	5.0%	
Inflation (i)	2.5%	
<i>Real Asset Returns:</i>		
	μ	σ
Stocks	6.4%	18.8%
Bonds	2.7%	9.2%
Money Market	0.7%	3.9%
<i>Asset Covariances:</i>		
Stocks-Bonds	0.4065%	
Bonds-Money Market	0.2033%	
Money Market-Stocks	0.0763%	

Table 6: Age Cutoff under Optimal Age-Based Default Rule and Change in Aggregate Default Wealth

	(1)	(2)	(3)	(4)
	$\alpha = 0$	$\alpha = 2$	$\alpha = 5$	$\alpha = 10$
Optimal age cutoff	44	47	36	20
G	+32.0%	+81.9%	+2.8%	+0.0%

Notes: The first row gives the integer value for the optimal cutoff age for different levels of risk aversion under the baseline assumptions (as shown in Table 5). The second row compares the aggregate default wealth generated under the optimal cutoff to the aggregate default wealth for a uniform DB default, or a cutoff age of 20, as specified in Equation (14).

Table 7: Characteristics of Employees Who Lose from the Optimal Age-Based Default Policy

	(1)	(2)	(3)	(4)
	$\alpha = 0$	$\alpha = 2$	$\alpha = 5$	$\alpha = 10$
Number defaulted into suboptimal plan	9	6	9	0
Number who are female	9	6	9	0
Number who have below-median income	7	6	9	0
Average loss in certainty equivalent	\$3,872	\$4,502	\$190	–
Average loss for women	\$3,872	\$4,502	\$190	–
Average loss for below-median income	\$3,902	\$4,502	\$190	–

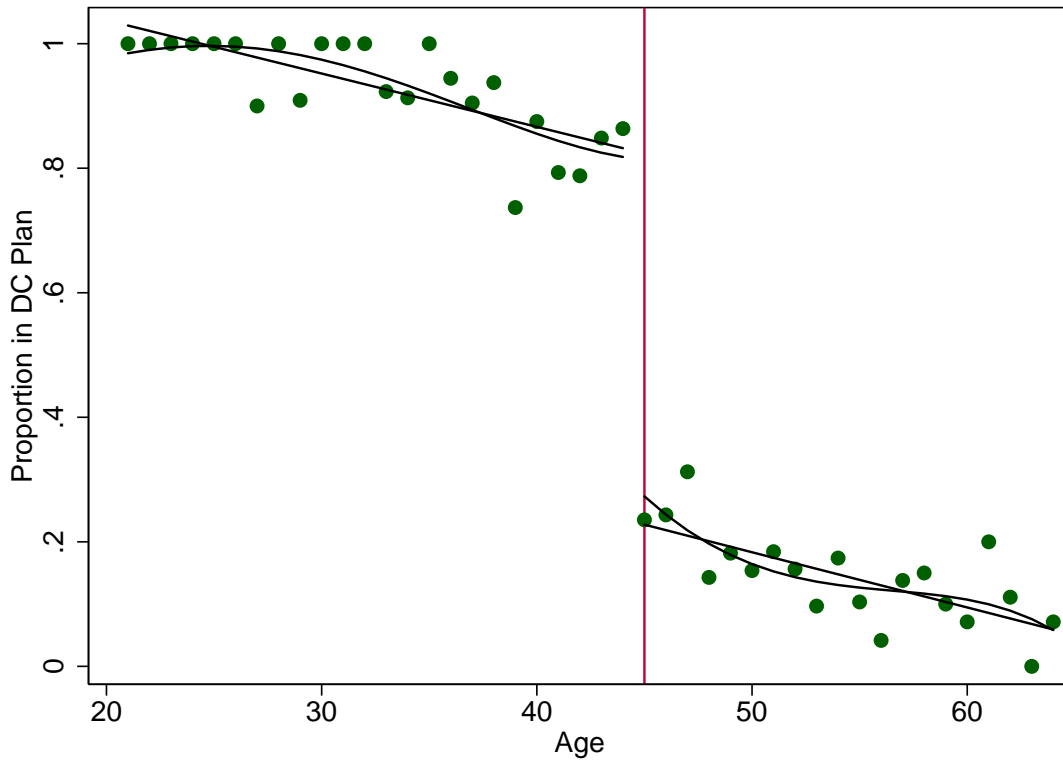
Notes: The first row summarizes the number of people for whom the default plan designated under the optimal cutoff policy has a lower certainty equivalent than the alternative. The second and third rows show how many of these individuals are women or workers with below-median income. The bottom panel summarizes the average reduction in the certainty equivalent for the individuals identified in the top panel.

Table 8: Comparative Statics for Optimal Age-Based Default Rule

Assumptions	$\alpha = 0$	$\alpha = 2$	$\alpha = 5$	$\alpha = 10$	$\bar{\alpha} = 2$
Baseline	44	47	36	20	47
5% DC Contribution Rate	36	40	20	20	40
10% DC Contribution Rate	47	49	42	20	49
1% DB Multiplier	56	57	56	30	57
3% DB Multiplier	38	42	20	20	41
100% Stocks	48	48	20	20	48
100% Bonds	32	39	33	20	39
100% Cash	20	26	23	20	20
0% Separation Hazard	40	36	23	20	36
10% Separation Hazard	46	48	36	20	48
No Investment Risk	40	47	50	50	48
Double Investment Risk	50	46	20	20	46
0% Real Wage Growth	45	47	36	20	47
4% Real Wage Growth	43	47	36	20	47
0% Real Discount Rate	44	47	36	20	47
2% Real Discount Rate	44	47	36	20	47
1.5% Inflation	42	45	20	20	45
3.5% Inflation	46	49	42	20	49

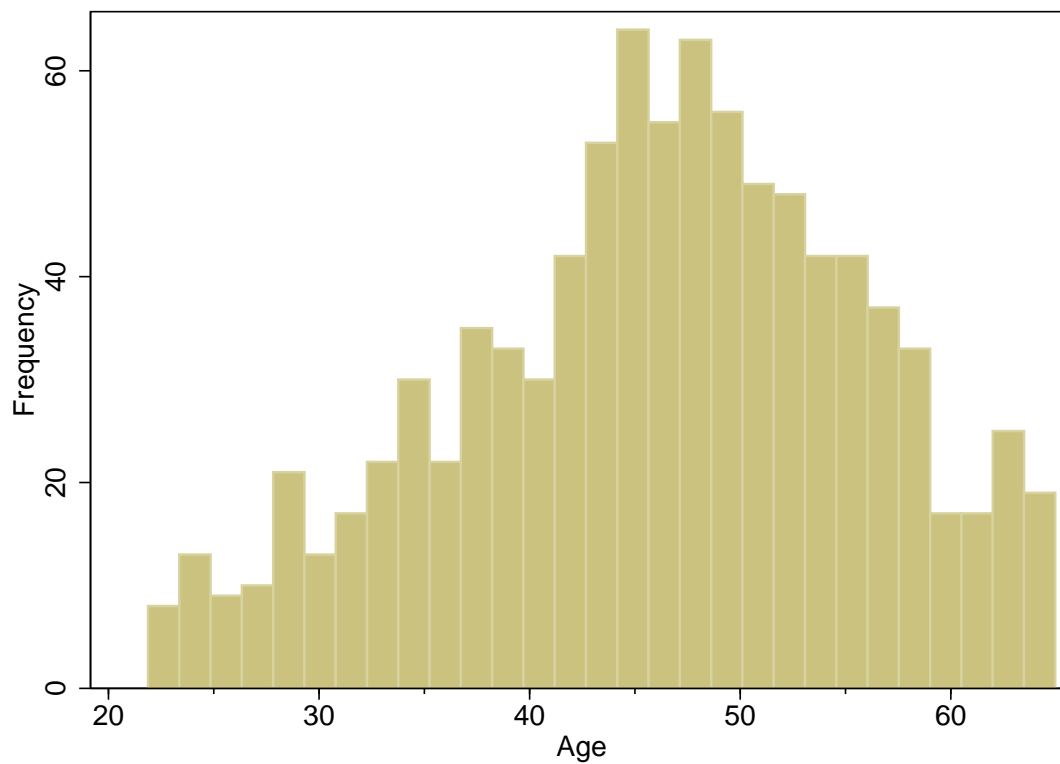
Notes: Each row gives the optimal cutoff age, a^* , for different levels of risk aversion for deviations from the baseline assumptions along one parameter dimension. Baseline assumptions are shown in Table 5. The last column assumes risk aversion is positively correlated with age but keeps the average level of risk aversion across all employees is equal to 2.

Figure 1: DC Enrollment Rate by Age



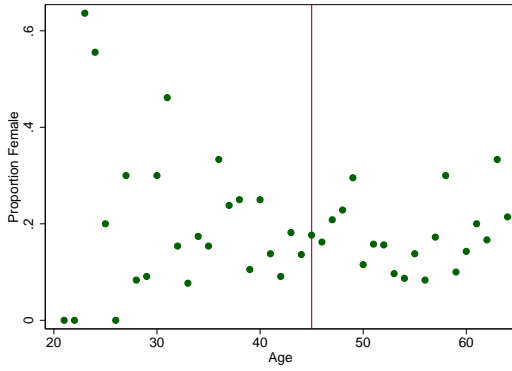
Notes: DC enrollment rates computed for each single year of age. Best fit line and cubic are shown for each side of the age 45 cutoff.

Figure 2: Distribution of Employee Age

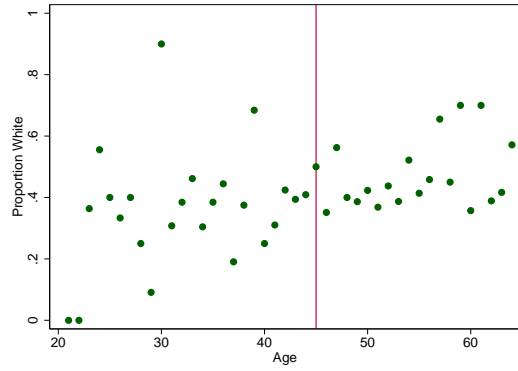


Notes: Histogram of employee age as of September 1, 2002, using one-year bins.

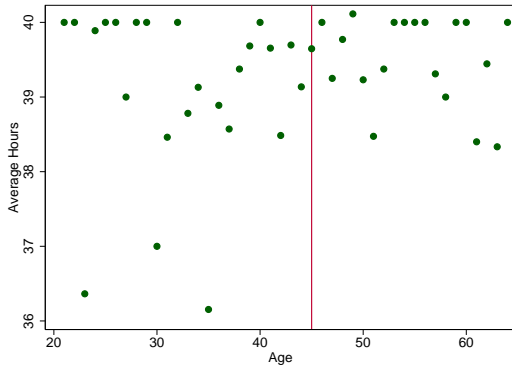
Figure 3: Average Value of Covariates by Single Year of Age



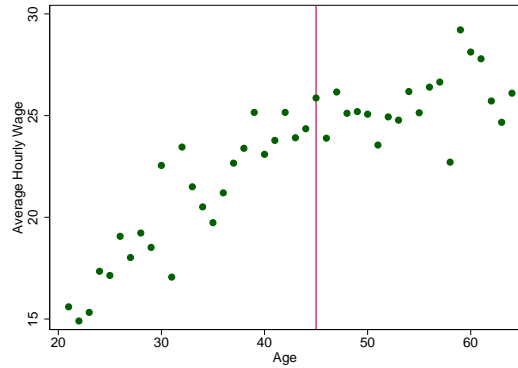
(a) Proportion Female



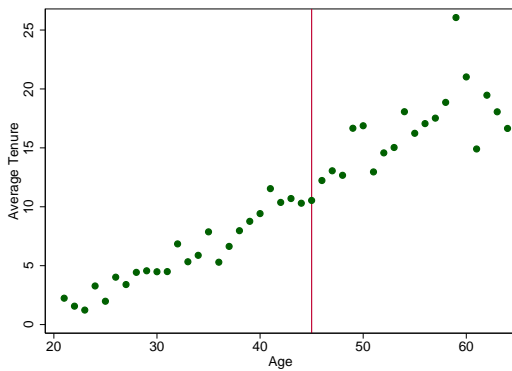
(b) Proportion White



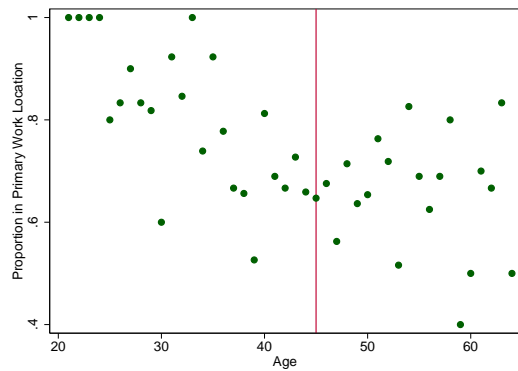
(c) Average Hours per Week



(d) Average Hourly Wage



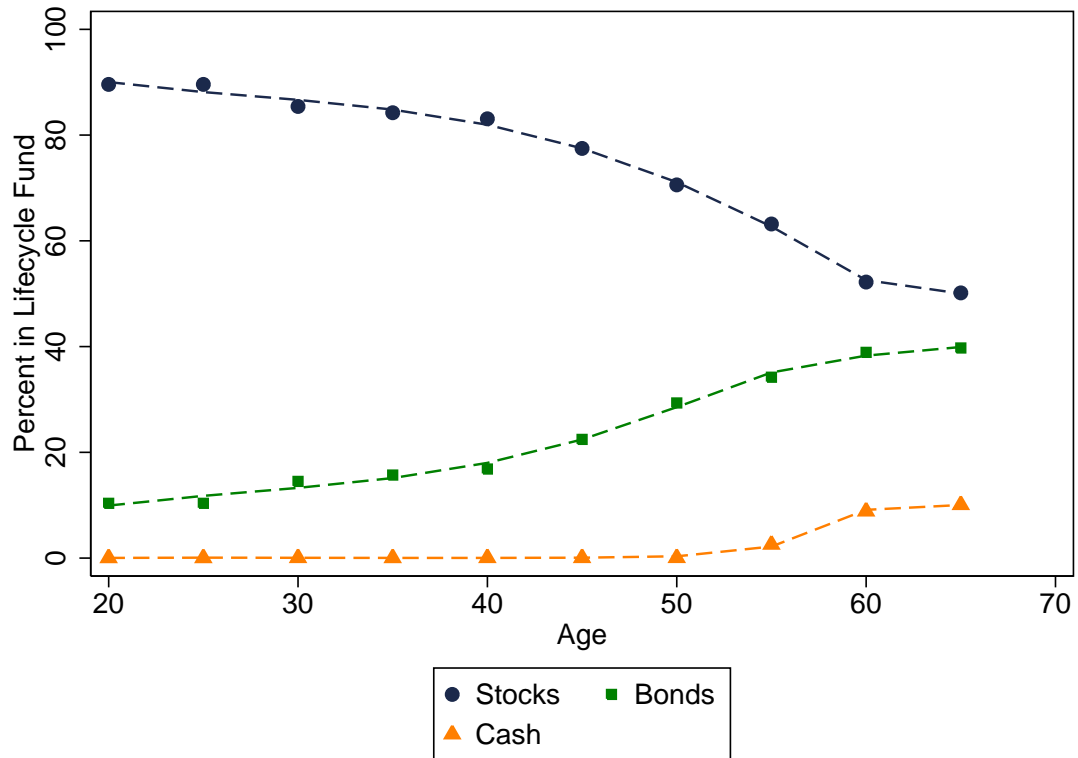
(e) Average Tenure



(f) Proportion in Primary Work Location

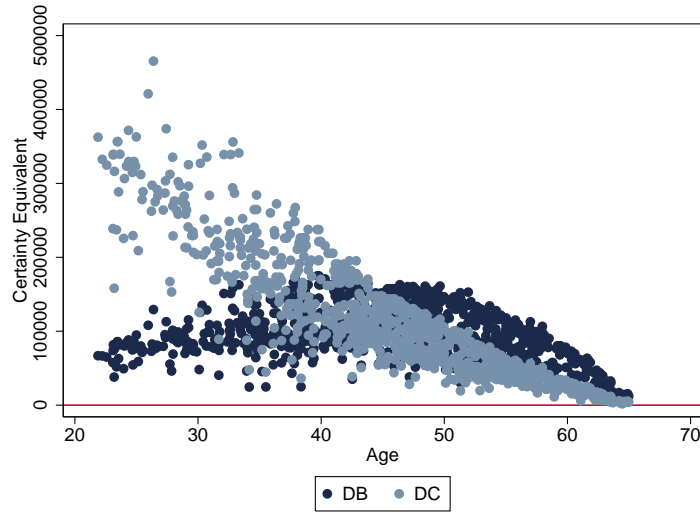
Notes: Panels used to verify no discontinuity at age 45 for other observable characteristics; vertical line marks age 45.

Figure 4: Asset Allocation by Age

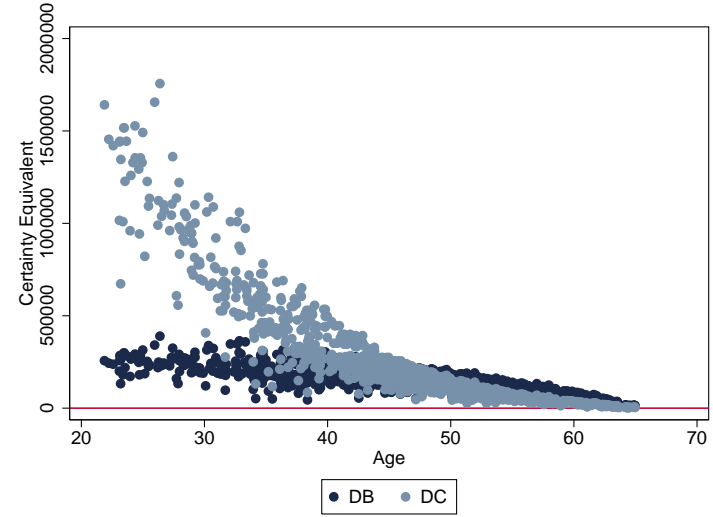


Notes: Markers represent asset allocation from Fidelity Freedom Funds (target-date retirement funds) for appropriate age assuming participant turns age 65 in target retirement year. Lines represent fit using fractional multinomial logit model with fourth-order age terms.

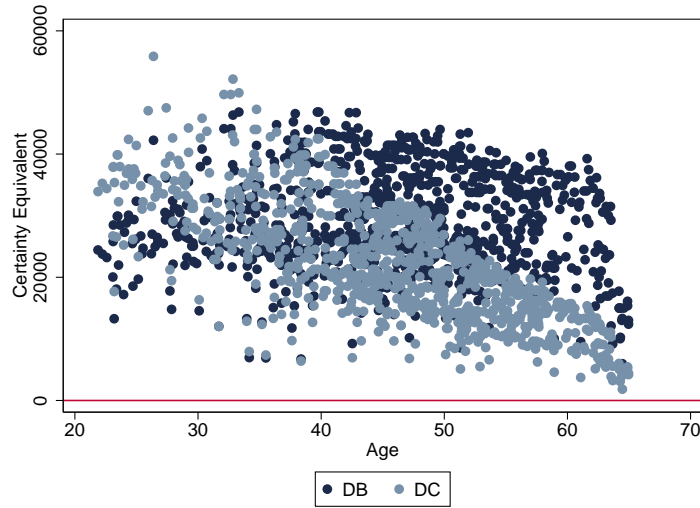
Figure 5: Certainty Equivalent Wealth by Age for Different Levels of Risk Aversion



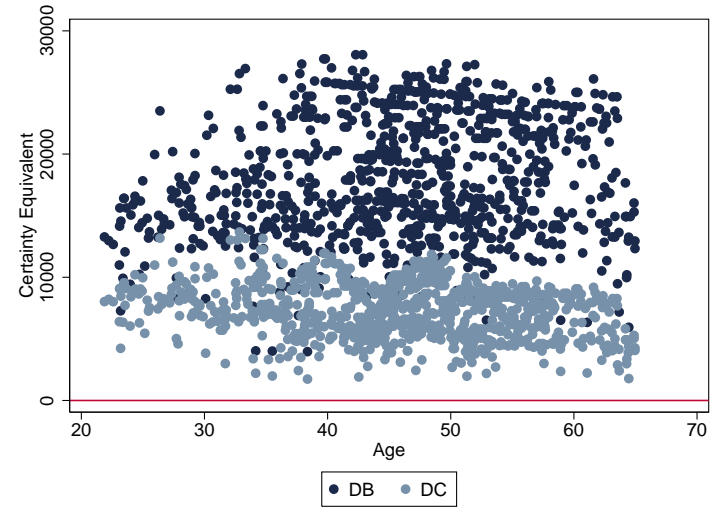
(a) $\alpha=0$ (*risk neutral*)



(b) $\alpha=2$



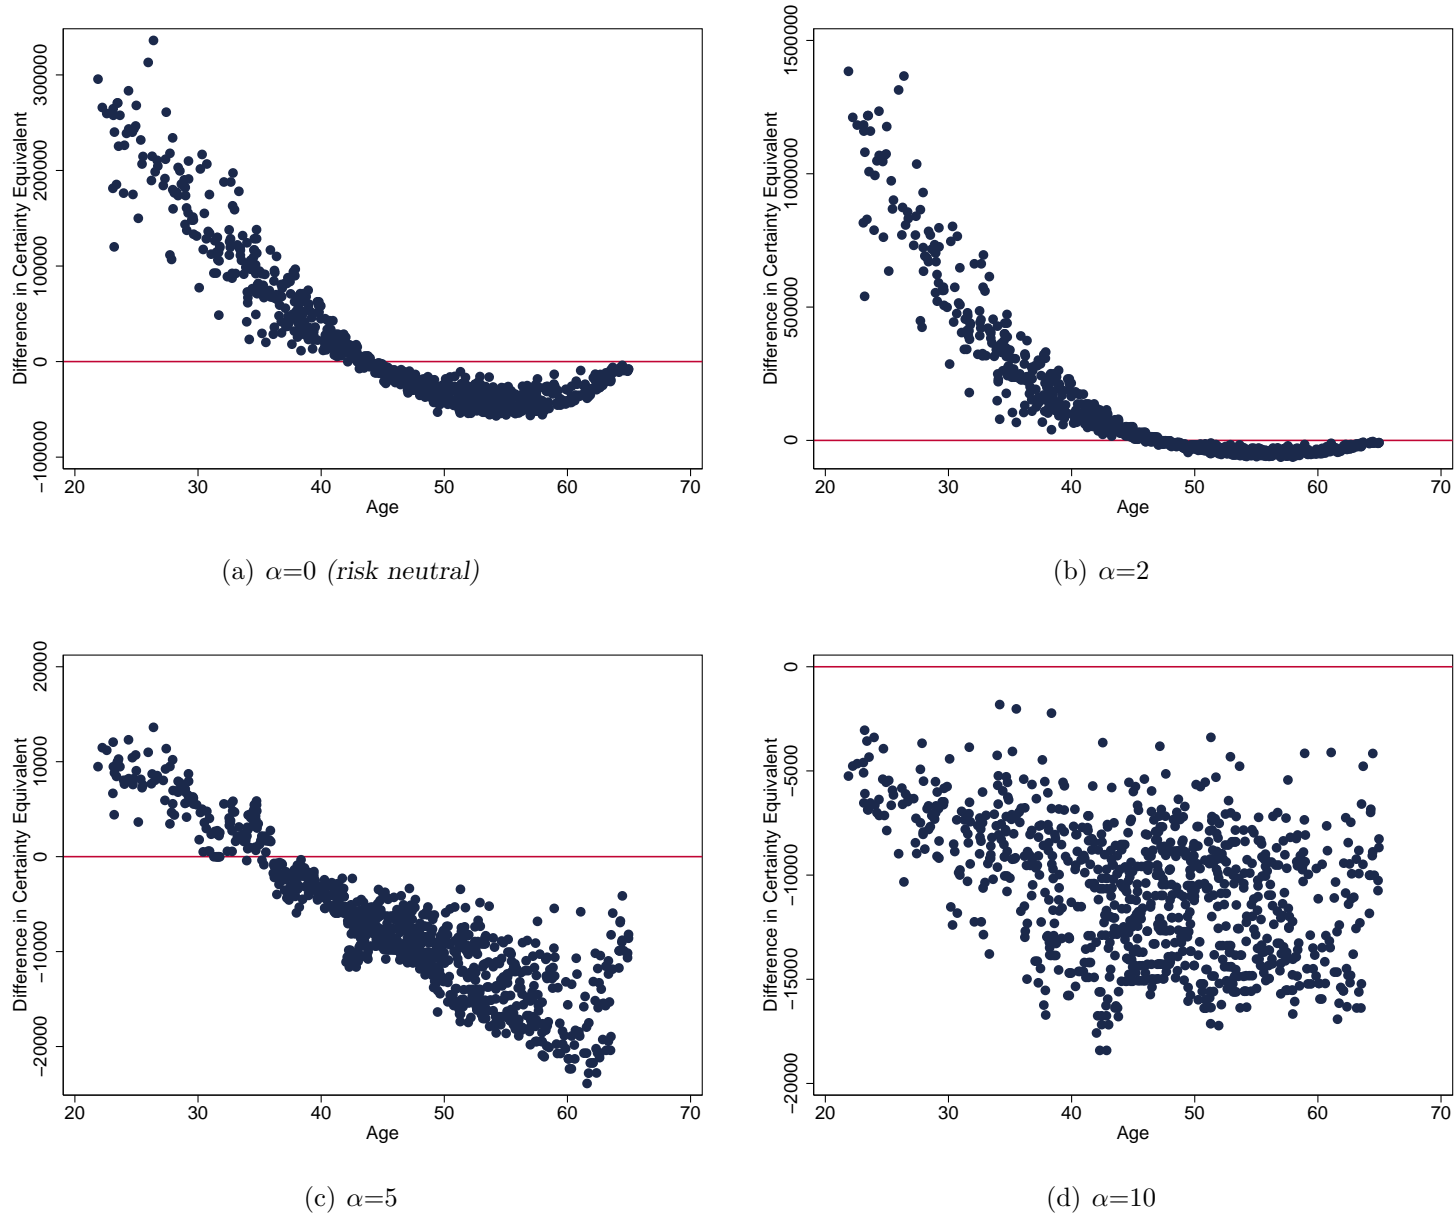
(c) $\alpha=5$



(d) $\alpha=10$

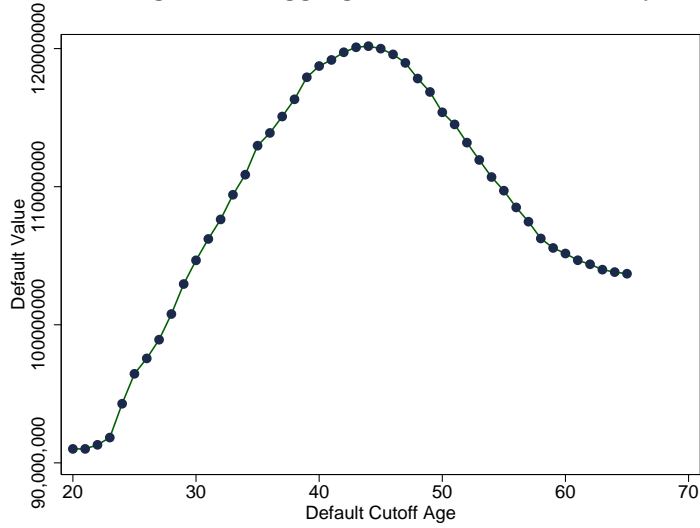
Notes: Each panel shows the certainty equivalent for each plan by age for different levels of risk aversion (α).

Figure 6: Difference in DB and DC Certainty Equivalent Wealth for Different Levels of Risk Aversion

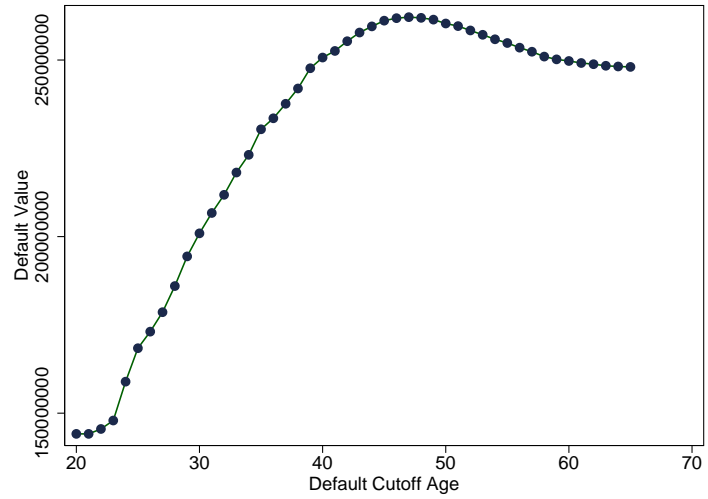


Notes: Each panel shows difference in certainty equivalent, $CE_i^{DC} - CE_i^{DB}$, for different levels of risk aversion (α).

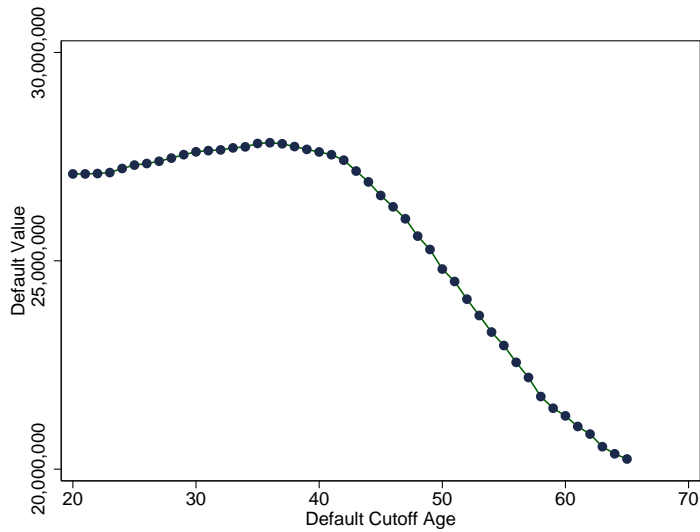
Figure 7: Aggregate Default Wealth by Cutoff Age for Different Levels of Risk Aversion



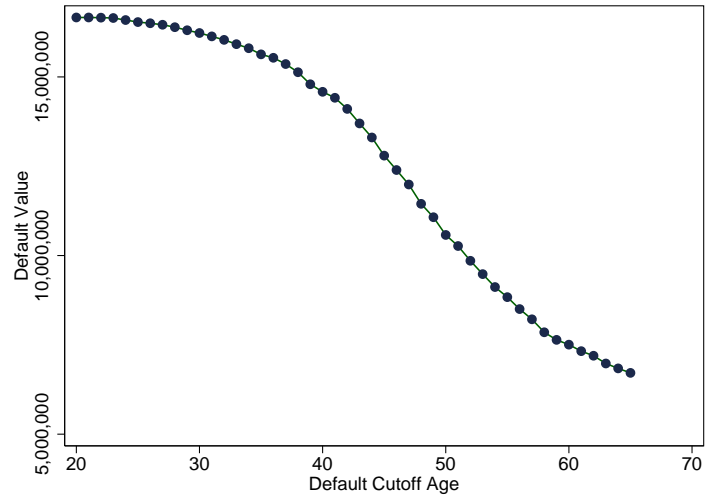
(a) $\alpha=0$ (*risk neutral*)



(b) $\alpha=2$



(c) $\alpha=5$



(d) $\alpha=10$

Notes: Each panel plots the aggregate default wealth by cutoff age for different levels of risk aversion (α). The optimal cutoff age maximizes the aggregate default wealth.

Appendix A: Robustness to Choice of Bandwidth

The regression discontinuity results shown in Tables 3 and 4 were done with a bandwidth of 5 years. The table below summarizes the regression discontinuity estimated for different values of the bandwidth across different specifications.

	(1)	(2)	(3)	(4)	(5)
LPM	0.584***	0.690***	0.703***	0.739***	0.737***
($h = 2.5$)	(0.058)	(0.108)	(0.112)	(0.196)	(0.200)
LPM	0.636***	0.544***	0.532***	0.697***	0.709***
($h = 7.5$)	(0.035)	(0.066)	(0.066)	(0.122)	(0.125)
LPM	0.665***	0.540***	0.538***	0.685***	0.694***
($h = 10$)	(0.030)	(0.059)	(0.059)	(0.108)	(0.110)
Probit	0.584***	0.687***	0.718***	0.728***	0.741***
($h = 2.5$)	(0.058)	(0.099)	(0.096)	(0.172)	(0.171)
Probit	0.636***	0.538***	0.542***	0.679***	0.708***
($h = 7.5$)	(0.035)	(0.071)	(0.074)	(0.114)	(0.114)
Probit	0.665***	0.527***	0.547***	0.683***	0.722***
($h = 10$)	(0.030)	(0.065)	(0.069)	(0.101)	(0.098)
Treatment of Age	Constant	Linear	Linear	Cubic	Cubic
Controls	No	No	Yes	No	Yes

Notes: Dependent variable is enrollment in the DC plan. Robust standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. LPM represents linear probability models. Probit models report marginal effects.

Appendix B: Comparative Statics for the Optimal Age-Based Default Policy

As shown in Equation (13) and restated below, the optimization problem and the implicit function theorem together provide a formula to sign the direction of changes in the optimal age cutoff for changes in known parameters:

$$\text{sign} \left(\frac{\partial a^*}{\partial \gamma_i} \right) = \text{sign} \left(\frac{\partial CE^{DC}}{\partial \gamma_i} - \frac{\partial CE^{DB}}{\partial \gamma_i} \right). \quad (\text{B-1})$$

Below, we examine several special cases to illustrate the intuition in this result.

Case 1: Suppose $c_j = c$ for all j . The optimal cutoff age a^* is increasing in c :

$$\text{sign} \left(\frac{\partial a^*}{\partial c} \right) = \text{sign} \left(\underbrace{\frac{\partial CE^{DC}}{\partial c}}_{>0} - \underbrace{\frac{\partial CE^{DB}}{\partial c}}_{=0} \right) > 0. \quad (\text{B-2})$$

Case 2: Similarly, suppose $b_j(w_j) = b$ for all j . The optimal cutoff age a^* is decreasing in b :

$$\text{sign} \left(\frac{\partial a^*}{\partial b} \right) = \text{sign} \left(\underbrace{\frac{\partial CE^{DC}}{\partial b}}_{=0} - \underbrace{\frac{\partial CE^{DB}}{\partial b}}_{>0} \right) < 0. \quad (\text{B-3})$$

Case 3: Suppose the utility function U is such that $U^{-1}(\beta w) = h(\beta) \cdot U^{-1}(w)$ for some function h . Then the maximization problem in Equation (10) does not depend on the discount rate d . Therefore:

$$\frac{\partial a^*}{\partial d} = 0. \quad (\text{B-4})$$

Appendix C: Monte Carlo Simulations

Monte Carlo simulations were performed to obtain 1,000 simulations of 45 years of asset returns. This methodology follows Shoven (1999). Suppose $r_i = \ln(1 + R_i)$, where R_i denotes the simple real return of one of three types of assets (stocks ($i = s$), bonds ($i = b$), and money market accounts ($i = m$)). We assume r_i is distributed normally, i.e. $1 + R_i$ is distributed lognormally. The lognormal distribution is skewed to the right and ensures that simple returns cannot fall below -100%. Let m_i , s_i , and s_{ij} denote the moments of R_i , and μ_i , σ_i , and σ_{ij} denote the moments of r_i .

The Monte Carlo simulation is done in three steps. First, μ_i , σ_i , and σ_{ij} are obtained for all classes of assets by the following set of equations:

$$\mu_i = \log \left(\frac{1 + m_i}{\sqrt{1 + \left(\frac{s_i}{1+m_i}\right)^2}} \right) \quad (\text{C-1})$$

$$\sigma_i^2 = \log \left(1 + \frac{s_i^2}{(1 + m_i)^2} \right) \quad (\text{C-2})$$

$$\sigma_{ij} = \log \left(1 + \frac{s_{ij}}{(1 + m_i)(1 + m_j)} \right) \quad (\text{C-3})$$

Next, three independent standard-normal random variables z are generated for each simulation using Matlab's random number generator. These three random numbers are combined such that the returns have the desired variances and covariances.

$$r_s = \mu_s + \sigma_s z_1 \quad (\text{C-4})$$

$$r_b = \mu_b + \frac{\sigma_{sb}}{\sigma_s} z_1 + z_2 \sqrt{\sigma_b^2 - \left(\frac{\sigma_{sb}}{\sigma_s}\right)^2} \quad (\text{C-5})$$

$$r_m = \mu_m + a z_1 + b z_2 + c z_3 \quad (\text{C-6})$$

The constants a , b , and c are given by:

$$a = \frac{\sigma_{sm}}{\sigma_s} \tag{C-7}$$

$$b = \frac{\sigma_{bm} - \frac{\sigma_{sm}\sigma_{sb}}{\sigma_s^2}}{\sqrt{\sigma_b^2 - \left(\frac{\sigma_{sb}}{\sigma_s}\right)^2}} \tag{C-8}$$

$$c = \sqrt{\sigma_m^2 - a^2 - b^2}. \tag{C-9}$$

Finally, the simple returns R are determined by using the transformation $R_i = \exp(r_i) - 1$.