Optimal Defaults and Active Decisions

Optimal Defaults and Active Decisions*

Gabriel D. Carroll James J. Choi David Laibson
Brigitte C. Madrian Andrew Metrick

November 11, 2008

Defaults often have a large influence on consumer decisions. We identify an overlooked but practical alternative to defaults: requiring individuals to make an explicit choice for themselves. We study such "active decisions" in the context of 401(k) saving. We find that compelling new hires to make active decisions about 401(k) enrollment raises the initial fraction that enroll by 28 percentage points relative to a standard opt-in enrollment procedure, producing a savings distribution three months after hire that would take 30 months to achieve under standard enrollment. We also present a model of 401(k) enrollment and derive conditions under which the optimal enrollment regime is automatic enrollment (i.e., default enrollment), standard enrollment (i.e., default non-enrollment), or active decisions (i.e., no default and compulsory choice). Active decisions are optimal when consumers have a strong propensity to procrastinate and savings preferences are highly heterogeneous. Financial illiteracy, however, favors default enrollment over active decision enrollment.

*We thank Hewitt Associates for their help in providing the data. We are particularly grateful to Lori Lucas, Jim McGhee, Scott Peterson, Deborah Bonus, and Yan Xu, some of our many contacts at Hewitt. We also appreciate many helpful conversations with Harlan Fleece and the exceptional research assistance of Nelson Uhan, Jared Gross, John Friedman, Hongyi Li, Laura Serban, Keith Ericson, Ben Grossman, and Fuad Faridi. Richard Thaler, Mark Iwry, Shlomo Benartzi, Annamaria Lusardi, and seminar participants at the University of Chicago, Harvard, MIT, Berkeley, Cornell, USC, LSE, Wharton, Dartmouth, Stanford, the New York Federal Reserve, the Russell Sage Foundation, and the NBER have provided much useful feedback. Carroll and Choi acknowledge financial support from National Science Foundation Graduate Research Fellowships. Choi acknowledges financial support from the Mustard Seed Foundation. Choi, Laibson, and Madrian acknowledge individual and collective financial support from the National Institute on Aging (grants R01-AG-16605, R29-AG-013020, and T32-AG00186). Laibson also acknowledges financial support from the Sloan Foundation. gdc@mit.edu, james.choi@yale.edu, dlaibson@harvard.edu, brigitte_madrian@harvard.edu, metrick@yale.edu.

Economists have studied two kinds of 401(k) enrollment. Under "standard enrollment," employees are by default not enrolled and can choose to opt into the plan. Under "automatic enrollment," employees are by default enrolled and can choose to opt out. In this paper, we analyze an overlooked third alternative: requiring employees to make an explicit choice for themselves before a deadline. In this "active decision" regime there is no default to fall back on because employees are not allowed to remain passive; they *must* explicitly declare their enrollment preference, regardless of what it is.

It might seem that defaults should not matter if agents believe they are arbitrarily chosen and if opting out of the default is easy. In practice, defaults tend to be sticky. For example, switching from a non-participation default to a participation default (automatic enrollment) can increase 401(k) participation rates among new hires by more than 50 percentage points, and about three-quarters of participants under automatic enrollment initially retain both the default contribution rate and the default asset allocation (Madrian and Shea 2001; Choi et al. 2002, 2004).[1] This perverse property of defaults has been documented in a wide range of other settings: organ donation decisions (Johnson and Goldstein 2003; Abadie and Gay 2004), car insurance plan choices (Johnson et al. 1993), car option purchases (Park, Yun, and MacInnis 2000), and consent to receive e-mail marketing (Johnson, Bellman, and Lohse 2002).

Because defaults powerfully influence outcomes without restricting choices, the intentional use of defaults as a policy lever has become increasingly common.[2] The Pension Protection Act of 2006 in the United States, the KiwiSaver Act in New Zealand, and the Pensions Act of 2007 in the United Kindom promote automatic enrollment in retirement

[1] Even bad defaults are sticky. In firms without automatic enrollment, Choi, Laibson, and Madrian (2007) show that many employees remain at the non-participation default even when this entails foregoing large arbitrage opportunities.

[2] Sunstein and Thaler (2003) coined the term "libertarian paternalism" to describe the approach of influencing behavior without restricting choices.

savings plans. In 2007, 26% of 401(k)-eligible U.S. employees worked for companies that automatically enrolled employees in their 401(k) plan (Harris Interactive 2007). Beshears, Choi, Laibson, Madrian, and Weller (forthcoming) summarize these policy developments in the retirement savings domain.

Defaults are socially desirable when a large majority of agents have a shared optimum and the default leads them to it (e.g., investing in a low-fee index fund). But even a well-chosen default may be undesirable if agents' optimal choices are highly heterogeneous. For example, in a firm whose workforce includes young, cash-strapped single parents and older employees who need to quickly build a retirement nest egg, one 401(k) savings rate isn't right for everyone. One could implement defaults that are tailored to each employee based on observable demographic characteristics, but unobserved employee heterogeneity may limit the helpfulness of such employee-specific defaults. In practice, the use of such differentiated defaults in the savings domain is limited, partly because of concerns about the legal ramifications of not treating all employees equally.[3]

Active decision mechanisms are an intriguing, though imperfect, alternative to defaults. On the positive side, active decisions avoid the biased outcomes introduced by defaults because active decisions do not corral agents into a uniform default choice. The active decision mechanism encourages agents to think about an important decision and avoid procrastinating. On the negative side, an active decision mechanism compels agents to struggle with a potentially time-consuming decision—which they may not be qualified to make—and then explicitly express their choice at a time which may be inconvenient. Some people would welcome a benign third party who is willing to make and automatically implement that decision for them. In addition, social engineers might prefer a default that aggressively encourages

[3] A notable exception is the use of target retirement funds as the default asset allocation in 401(k) plans with automatic enrollment. Such funds, which were formally sanctioned in regulations issued by the Department of Labor pursuant to the implementation of the Pension Protection Act of 2006, are now the default in half of plans with automatic enrollment (Hewitt 2007).

some social goal, like organ donation or retirement saving.[4]

The current paper lays the groundwork for a debate about active decisions by describing a natural experiment at one large firm. This firm (unintentionally) used an active decision 401(k) enrollment regime which required employees to explicitly state an enrollment preference within 30 days of hire. There was no formal penalty for failing to actively express an enrollment preference. In fact, there was an unstated non-enrollment default; employees who refused to declare a preference were not enrolled in the 401(k). Nonetheless, compliance with the active decision requirement was nearly universal.

The firm then switched to a standard enrollment regime as a by-product of the transition from a paper-based administrative system to a phone-based administrative system. The firm did not anticipate that the transition to a phone-based system with a default of non-enrollment would transform the psychology of 401(k) participation. Rather, the change in administrative systems was motivated solely by the convenience and efficiency of phone-based enrollment. The loss of active decision effects was a collateral consequence of that transition.

We find that active decisions raised the initial fraction of employees enrolled by 28 percentage points relative to what was obtained with a standard default of non-enrollment. Active decisions raised average savings rates and accumulated balances by accelerating decision-making. We show that conditional on demographics, each employee under active decisions on average *immediately* chose a savings rate similar to what she would otherwise have taken up to thirty months to attain under standard enrollment. Because the typical worker will change jobs several times before retirement, accelerating the 401(k) savings decision by more than two years at the beginning of each job transition could have a large impact on accumulated wealth at retirement.

[4] There is a large literature debating whether U.S. households are undersaving for retirement. See Skinner (2007) for a survey.

Given that enrollment mechanism design choices cannot be avoided and have large effects on outcomes (Sunstein and Thaler 2003), how should a benign planner decide which mechanism to implement? We present a model that provides a formal framework for evaluating the welfare consequences of different enrollment mechanisms.

In our model, defaults matter for two reasons. First, the process of opting out of the default may be more or less costly at different times; this creates option value to waiting for a low-cost time to take action. Second, agents with present-biased preferences (Phelps and Pollak 1968; Laibson 1997; O'Donoghue and Rabin 1999b) may procrastinate in their decision to opt out of the default. Taking individuals' utility without present bias as normative, we derive conditions under which a benign utilitarian social planner should implement automatic enrollment, standard enrollment, or active decisions. Active decisions are socially optimal when consumers have highly heterogeneous optimal savings rates and a strong propensity to procrastinate. Heterogeneity may arise from a host of factors including past and future receipt of bequests, spousal savings, household variation in historical portfolio returns, and lifecycle liquidity needs (e.g., children).

The rest of this paper has the following structure. Section 1 describes the details of the two 401(k) enrollment regimes at the company we study. Section 2 describes our data. Section 3 compares the 401(k) savings decisions of employees hired under the active decision regime to those hired under the standard enrollment regime. Section 4 presents a model of procrastination for time-inconsistent agents with rational expectations and uses this model to derive the socially (second-best) optimal enrollment mechanism for such agents. Section 5 discusses the key implications of the model, as well as unmodeled factors that affect 401(k) choices. In particular, defaults will be better-suited for domains where individual agents have worse information about their optima than the social planner. Section 6 concludes and briefly discusses practical considerations in the implementation of active decision mechanisms.

1 The Natural Experiment

We use employee-level data from a publicly traded Fortune 500 company in the financial services industry. In December 1999, this firm had offices in all fifty states, as well as the District of Columbia and Puerto Rico. This paper will consider the 401(k) savings decisions of employees at the firm from January 1997 through December 2001.

Until November 1997, all newly hired full-time employees at the firm were *required* to submit a form within 30 days of their hire date stating their 401(k) participation preferences, *regardless* of whether they wished to enroll or not. Although there was no tangible penalty for not returning the 401(k) form, human resource officers report that only 5% of employees failed to do so. We believe that this high compliance rate arose because the form was part of a packet that included other forms whose submission was legally required (e.g., employment eligibility verification forms and tax withholding forms). Moreover, employees who did not return the form were reminded to do so by the human resources department. A failure to return the form was treated as a decision not to enroll in the 401(k), although this non-enrollment default was not publicized.

Employees who declined to participate in the 401(k) plan during this initial enrollment period could not subsequently enroll in the plan until the beginning (January 1) of succeeding calendar years. Later in the paper, we will show that this delay did not drive the active decision effect.

At the beginning of November 1997, the company switched from a paper-based 401(k) enrollment system to a telephone-based system. Employees hired after this change no longer received the active decision 401(k) enrollment form when hired. Instead, they were given a toll-free phone number to call if and when they wished to enroll in the 401(k) plan. We call this new system the "standard enrollment" regime because its non-enrollment default is used by most companies. The telephone-based system also allowed employees to enroll on a

daily basis, rather than only at the beginning of each calendar year as had previously been the case. This change applied not only to employees hired after November 1997, but to all employees working at the company.[5]

A number of other 401(k) plan features also changed at the same time. We believe that these additional changes made 401(k) participation more attractive, so our estimates of the active decision effect are a lower bound on the true effect. These other changes include a switch from monthly to daily account valuation, the introduction of 401(k) loans, the addition of two new funds as well as employer stock to the 401(k) investment portfolio,[6] and a switch from annual to quarterly 401(k) statements. Table I summarizes the 401(k) plan rules before and after the November 1997 changes.

2 The Data

We have two types of employee data. The first dataset is a series of cross-sections at year-ends 1998, 1999, 2000, and 2001. Each cross-section contains demographic information for everybody employed by the company at the time, including birth date, hire date, gender, marital status, state of residence, and salary. For 401(k) plan participants, each cross-section also contains the date of enrollment and year-end information on balances, asset allocation, and the terms of any outstanding 401(k) loans. The second dataset is a longitudinal history of every individual transaction in the plan from September 1997 through April 2002: savings rate elections, asset allocation elections for contributions, trades among funds, loan-based withdrawals and repayments, financial hardship withdrawals, retirement withdrawals, and

[5] Unfortunately, we have no data with which to separately identify an active decision effect from a paperless enrollment effect. We have been unable to find any empirical studies documenting whether 401(k) participation is materially affected by changing from a paper-based enrollment scheme to an electronic enrollment technology.

[6] Prior to November 1997, employer stock was available as an investment option only for match balances and contributions made with after-tax money.

rollovers.

To analyze the impact of active decisions, we compare the behavior of two groups: employees hired between January 1, 1997 and July 31, 1997 under the active decision regime, and employees hired between January 1, 1998 and July 31, 1998 under the standard enrollment regime. We refer to the first group as the "active decision cohort" and the second group as the "standard enrollment cohort."

We exclude employees hired prior to January 1, 1997 because the company made two plan changes that took effect on January 1, 1997. First, the company eliminated a one-year service requirement for 401(k) eligibility. Second, the company changed the structure of its 401(k) match. Although active decisions were used until the end of October 1997, we do not include employees hired from August through October to avoid any confounds produced by the transition to standard enrollment. For example, an enrollment blackout was implemented for several weeks during the transition.

Our key identifying assumptions are that the two cohorts have similar savings preferences on average, and that any common shocks unrelated to the enrollment mechanism which affected 401(k) contribution behavior are small. Company officials reported no material changes in hiring or employment practices until shortly before the 2001 recession. In addition, the economic environment faced by these two groups of employees early in their tenure was similar.

The active decision cohort is first observed in our cross-sectional data in December 1998 (seventeen to 24 months after hire) and in the longitudinal data starting at the end of September 1997 (two to nine months after hire). The longitudinal data are not sufficient for calculating participation rates because they do not include 401(k) non-participants. The cross-sectional data contain both participants (along with their enrollment dates) and non-participants, but they do not contain employees who left the firm before the snapshot date.

Therefore, any participation rate calculation for the active decision cohort must be among the subset that eventually attains at least seventeen to 24 months of tenure at the company.

The standard enrollment cohort is also observed in our cross-sectional data starting in December 1998, which is only five to twelve months after their hire date. Therefore, we can calculate participation rates for the standard enrollment cohort among the subset that eventually attains only five to twelve months of tenure. In the longitudinal data, 401(k) participants from this cohort are observed as soon as they enroll.

Even if active decisions had no effect on participation, the data introduces differential selection effects that could falsely inflate participation rates calculated for the active decision cohort relative to the standard enrollment cohort. For example, if there are fixed enrollment costs and 401(k) benefits accrue over time at the job, then people who know they will stay at the company longer are more likely to join the 401(k). All else equal, a population that attains seventeen to 24 months of tenure will have had a higher participation rate shortly after hire than a population that attains only five to twelve months of tenure.[7]

To equalize the sample selectivity of both the active decision and standard enrollment cohorts, we drop standard enrollment cohort employees who were not still at the company at year-end 1999. Thus, presence in both samples is conditional on staying at the company at least seventeen to 24 months. Our results are not qualitatively affected if we do not impose this equalizing sample restriction on the standard enrollment cohort.

Table II presents demographic statistics on the active decision and standard enrollment cohorts at the end of December in the year *after* they were hired. The cohorts are similar in age, gender composition, income, and geographical distribution. The dimension along which

[7]Even and MacPherson (2005) find that under standard enrollment, employees who enroll in their 401(k) are less likely to leave the company than non-participants. This correlation, however, does not imply that 401(k) participation *causes* lower attrition. In unpublished research, we find no attrition rate discontinuities between cohorts hired immediately before and immediately after transitions between standard enrollment and automatic enrollment, regime shifts that produce much larger participation changes than those documented in the current paper.

they differ most is marital status, and even here the differences are not large: 56% of the active decision cohort is married, while this is true for only 50% of the standard enrollment cohort.[8] The third column of Table II shows that the new-hire cohorts are different from employees at the company overall. As expected, the new-hire cohorts are younger, less likely to be married, and paid less on average. The last column reports statistics from the Current Population Survey, providing a comparison between the company's employees and the total U.S. workforce. The company has a relatively high fraction of female employees, probably because it is in the service sector. Employees at the company also have relatively high salaries. This is partially due to the fact that the company does not employ a representative fraction of very young employees, who are more likely to work part-time and at lower wages.

3 Empirical Results

3.1 401(k) Enrollment

We first examine the impact of active decisions on 401(k) participation. Figure I plots the fraction enrolled in the 401(k) after three months of tenure for employees who were hired in the first seven months of 1997 (the active decision cohort) and the first seven months of 1998 (the standard enrollment cohort). We use the third month of tenure because it could take up to three months for enrollments to be processed in the active decision regime.[9] The average three-month enrollment rate is 69% for the active decision cohort, versus 41% for the standard enrollment cohort, and this difference is statistically significant at the 1% level

[8] Because we are able to estimate these summary statistics precisely, the differences between the two cohorts are statistically significant at the 1% level along all the dimensions shown in Table II. Adjusting for these differences through regressions leaves our results qualitatively unchanged.

[9] Enrollments were only processed on the first of each month under the active decision regime. Since employees had thirty days to turn in their form, an employee who was hired late in a month and turned in her form just before the deadline could be enrolled in the third month after her hire.

for every hire month.[10]

Figure II plots the fraction of employees who have enrolled in the 401(k) plan against tenure. The active decision cohort's enrollment rate grows more slowly than the standard enrollment cohort's, so the enrollment gap decreases with tenure. Nonetheless, the active decision cohort's enrollment rate exceeds the standard enrollment cohort's by 17 percentage points at 24 months of tenure, and by 5 percentage points at 42 months. These differences are statistically significant at the 1% level for every tenure level after the first month.

Figures I and II could be misleading if enrollees under the active decision regime are subsequently more likely to stop contributing to the 401(k) plan. However, attrition rates from the 401(k) plan are indistinguishable under the active decision regime and the standard enrollment regime. Indeed, 401(k) participation is a nearly absorbing state under either enrollment regime.[11]

We ascribe the active decision effect to the fact that active decision employees had to express their 401(k) participation decision during their first month of employment, rather than being able to delay action indefinitely. However, there is another distinction between the active decision and standard enrollment regimes as implemented at the company. Under the standard enrollment regime, employees could enroll in the 401(k) plan at any time. Under the active decision regime, if employees did not enroll in the plan in their first thirty days at the company, their next enrollment opportunity did not come until January 1 of the following calendar year.[12] Therefore, in addition to the required affirmative or negative

[10] Our analysis focuses on employees hired from January to July 1997 and January to July 1998 in order to control for seasonalities in hiring. However, employees hired during the first two months of standard enrollment—November and December 1997—have three-month participation rates very similar to the January to July 1998 standard enrollment hires: 43% (November) and 46% (December).

[11] These calculations are available from the authors.

[12] In fact, the active decision cohort we analyze (January to June 1997 hires) was able to enroll in the 401(k) plan any time after November 1997, when the company switched to the phone-based daily enrollment system. At hire, however, the active decision employees were not aware of this impending change and would have believed January 1, 1998 to be their next enrollment opportunity.

enrollment decision, the active decision cohort faced a narrower enrollment window than the standard enrollment cohort. In theory, this limited enrollment window could cause higher initial 401(k) enrollment rates through a simple discreteness effect: employees who would have otherwise enrolled between their second month of tenure and the following January instead enroll in their first thirty days.

However, the enrollment differences between the cohorts are too large to be explained by discreteness alone. If only the discreteness effect were operative, participation rates for the two groups should be equal after twelve months of tenure. In fact, the participation rate of the active decision cohort at three months of tenure is not reached by the standard enrollment cohort until thirty months of tenure.

3.2 401(k) Contribution Rate

Although active decisions induce earlier 401(k) enrollment, this may come at the cost of more careful and deliberate thinking about how much to save for retirement. We now turn our focus to the impact of active decisions on other 401(k) savings outcomes.

Figure III plots the relationship between tenure and the average 401(k) contribution rate for the active decision and standard enrollment cohorts. These averages include both participants (who have a non-zero contribution rate) and non-participants (who have a zero contribution rate). Because our longitudinal data do not start until September 1997, the contribution rate profile cannot be computed for the entire active decision cohort until nine months of tenure.

The active decision cohort contributes 4.8% of income on average at month nine, and this slowly increases to 5.5% of income by the fourth year of employment. In contrast, the standard enrollment cohort contributes only 3.6% of income on average at month nine, and it takes more than 33 months for it to reach the active decision cohort's nine-month

savings rate. At each tenure level in the graph, the difference between the groups' average contribution rates is statistically significant at the 1% level.

Figure IV plots the average contribution rate of employees who have a non-zero contribution rate (i.e., 401(k) participants). In contrast to Figure III, active decision participants have a lower average contribution rate than standard enrollment participants until the fourth year of tenure.[13] Examining the distribution of contribution rates in each cohort (not shown), we see that at each percentile, the active decision cohort's contribution rate matches or exceeds the standard enrollment cohort's at virtually every tenure level. Most of the difference between the two cohorts' distributions is due to the active decision cohort employees signing up for the 401(k) earlier in their tenure. Therefore, the lower average contribution rate among active decision participants is not due to active decisions lowering the savings rates of those who would have otherwise contributed more under standard enrollment. Rather, active decisions bring employees with weaker savings motives into the participant pool earlier in their tenure.

Table III presents the results of a tobit regression of the two regimes' contribution rates on demographic variables. The contribution rate is censored below at 0% and above at 17% of pay, reflecting the plan's contribution limits. Both active decision and standard enrollment employees are included in the regression, regardless of participation status. If the employee was hired under the standard enrollment regime, the dependent variable is equal to the contribution rate at thirty months after hire. If the employee was hired under the active decision regime, the dependent variable is equal to an estimate of the contribution rate at three months after hire. This estimate is constructed by taking the earliest available contribution rate (which may be as late as nine months after hire) for the active decision employee and setting that contribution rate to zero if the employee had not enrolled in the

[13] These differences are statistically different at the 1% level through the 29th month of tenure, and at the 5% level through the thirtieth month of tenure.

plan within three months of hire. The explanatory variables are a constant, log of salary, and gender, marital status, and age dummies. The effect of these variables is allowed to vary between the active decision and standard enrollment cohorts. To test the hypothesis that savings rates are more haphazard under active decisions, we also allow the variance of the error term to vary across the two cohorts.[14]

The regression coefficients suggest that in expectation, there is little difference between the savings rate an employee chooses immediately after hire under active decisions and the rate she would have in effect thirty months after hire under standard enrollment. The only variable we can statistically reject having the same effect under both regimes is gender; at our company, women save somewhat less than men under active decisions but not under standard enrollment. The error term in the regression is significantly smaller for the active decision cohort than for the standard enrollment cohort, suggesting that the rush of the active decision deadline does not cause people to make more haphazard savings rate decisions.[15]

In sum, active decisions cause employees to immediately choose a savings rate that is similar on average to what they would take up to thirty months to attain under standard enrollment.[16]

3.3 401(k) Asset Allocation

The effect of active decisions on asset allocation outcomes cannot be cleanly inferred from the natural experiment we study here because the menu of investment fund options changed

[14] To equalize sample selectivity for the two cohorts, we restrict both samples to employees who remain in our data for 30 months. This is why the number of data points in the regression is less than the total number of employees in the two cohorts.

[15] On the other hand, if idiosyncratic variation in the savings rate is primarily due to employees optimally utilizing private information, then a lower error term variance could be consistent with a less well-considered savings decision.

[16] In unreported results, we find that among standard enrollment employees who stay at the company for at least thirty months and who enroll within those thirty months, those who are married, older, and have higher salaries enroll sooner.

in November 1997, the same time that the company switched from active decisions to the standard enrollment regime. Prior to the change, employer stock was only available as an investment option for after-tax contributions, and few employees made after-tax (as opposed to pre-tax) contributions.[17] During the transition to standard enrollment, employer stock was added as an investment option for pre-tax 401(k) contributions. Subsequently, the average allocation to employer stock more than doubled and the average allocation to all other asset classes correspondingly decreased. It is impossible to determine how much of this increase was caused by the standard enrollment regime, and how much was caused by the dramatic increase in the fraction of employees for whom employer stock was a viable investment option.

The impact of active decisions on asset allocation is an important open question, since many individuals have low levels of financial knowledge about different asset classes (John Hancock 2002) and tend to make poor asset allocation choices (Benartzi and Thaler 2001; Cronqvist and Thaler 2004; Choi, Laibson, and Madrian 2008). In section 5.2, we explain why active decisions are likely to be better suited for contribution rate choices than for asset allocation choices. Asset allocation decisions are probably best treated with a clear default option.

3.4 401(k) Asset Accumulation

We next consider the impact of active decisions on asset accumulation. Asset accumulation analysis is confounded by time effects, since asset returns are highly volatile. Moreover, the investment fund menu changed over time, as explained above, further confounding this analysis. Nonetheless, it is the level of asset accumulation that will ultimately drive retire-

[17] Pre-tax contributions are more tax-efficient unless the contributor has a short investment horizon and expects tax rates to rise sharply in the future. At the company, 14% of 401(k) participants made after-tax contributions during 1998.

ment timing and consumption levels. Studying asset accumulation also gives us insight into whether increased contribution rates under active decisions are offset by increased 401(k) loan activity and withdrawals.[18]

To measure asset accumulation, we divide 401(k) balances by annual base pay. Our measure of 401(k) balances excludes outstanding principal from 401(k) loans and any balances an employee rolled over from a previous employer.

Figure V reports balance-to-pay ratios at the 25th, 50th, 75th, and 90th percentiles of the balance-to-pay distribution for the active decision and standard enrollment cohorts. The impact of the market downturn in 2001 appears around the 48th month of tenure for the active decision cohort and the 36th month of tenure for the standard enrollment cohort.

It is apparent that the balance-to-pay ratio paths are nearly identical for the two cohorts at both the 75th and 90th percentiles. In contrast, the 25th percentile active decision employee has a much higher balance-to-pay ratio because participation begins two years earlier in her tenure than it does for the 25th percentile standard enrollment employee. The 50th percentile active decision employee has a slightly higher balance-to-pay ratio, but the effects of the 2001 stock market downturn muddy the picture. Overall, it appears that active decision enrollment only affects asset accumulation in the bottom half of the accumulation distribution.[19]

Unfortunately, we do not have any data on employees' assets outside of the 401(k). An important area for future research is the extent to which interventions like active decision and automatic enrollment affect total wealth accumulation, versus simply reshuffling assets between accounts. In the case of the active decision company studied here, it seems likely

[18] The active decision cohort did not have 401(k) loans available to them at the time they made their initial contribution rate decision. However, after November 1997, they were able to borrow against their 401(k) balances.

[19] We also find that the active decision cohort is more likely than the standard enrollment cohort to take withdrawals and loans from their 401(k). This is unsurprising, since the active decision cohort has higher 401(k) balances, and one must have 401(k) balances in order to withdraw or borrow against them.

that the total wealth impact was positive, given the extremely high employer match rate on 401(k) contributions (see Table I).

4 A Model of Optimal Enrollment Regimes

The empirical analysis in Section 3 shows that active decisions accelerate 401(k) enrollment. But these results do not enable us to evaluate the welfare consequences of active decisions. We now present a model that provides a framework for thinking about socially optimal enrollment regimes when agents possibly suffer from present bias.[20] An earlier draft of this paper (Carroll et al. 2007) presents results for the model under more general distributional/preference assumptions, as well as results showing that active decision regimes are particularly good enrollment mechanisms for agents who are naive about their future tendency to procrastinate. Proofs of all results are in the online appendix.

4.1 The Employee's Problem

Consider an employee who is initially at a default 401(k) savings rate d. Each period she draws a stochastic time cost c and decides whether to pay this time cost to opt out of the default and move to her optimal 401(k) savings rate s. If the employee opts out, she immediately incurs cost c, and she suffers no future losses because she is now at her optimum. If the employee does not opt out, she faces a flow utility loss $L = \kappa(s-d)^2$ at the beginning of the next period (κ is an exogenous constant), and the process repeats.[21] We assume that c is uniformly distributed on the interval $[\underline{c}, \bar{c}]$ (where $0 < \underline{c} < \bar{c}$) and independent across

[20] The theoretical analysis in Choi et al. (2003) is based upon the model presented below, but uses a two state cost distribution and does not derive conditions under which active decision regimes are optimal.

[21] In firms that match employee 401(k) contributions up to a threshold, the utility loss function may be kinked at this threshold (which may or may not coincide with the worker's optimum). In this case, our analytically convenient loss function is unrealistic. Loss aversion—which we do not model here—may also affect the loss function.

periods. We also assume for simplicity that s is constant over time.

The employee has quasi-hyperbolic preferences (Phelps and Pollak 1968; Laibson 1997): she has a long-term discount factor δ and an *additional* short-term discount factor β, where $\beta, \delta \in (0, 1]$. Thus, if her utility in periods $0, 1, 2, \ldots$ is $u_0, u_1, u_2, \ldots$, then her intertemporal utility from the perspective of self t is $U_t = u_t + \beta(\delta u_{t+1} + \delta^2 u_{t+2} + \cdots)$. If $\beta < 1$, the agent suffers from dynamic inconsistency. The resulting decision problem can be modeled as a game among the different selves. We assume that the employee is a sophisticated quasi-hyperbolic discounter, meaning that the values of β and δ are common knowledge among all the selves.

The employee will take action if and only if she draws a cost c less than some cutoff cost c^*. We restrict our attention to stationary equilibria, where the same c^* is used each period. As we show in the online appendix, c^* is uniquely determined and weakly increasing in L and β. Under active decisions, the employee is compelled to always opt out in the first period.

The assumption of time-invariant s is consistent with Table III, which shows that the contribution rate employees choose under active decisions is similar to the contribution rate they would have eventually chosen under standard enrollment. The assumption $L = \kappa(s-d)^2$ implies that c^* is an increasing function of $|s - d|$, which means that outside of an active decision regime, the expected delay until opt-out decreases with $|s - d|$. The data support this prediction. In the standard enrollment cohort, the default contribution rate d is 0, and s is the contribution rate initially chosen when an employee opts out of the default and starts contributing to the plan. The graph in Figure VI shows that people who choose a higher initial contribution rate (and hence have a higher $|s - d|$) on average opt out of the non-enrollment default sooner.[22]

[22] In unpublished research, we have verified in another company with automatic enrollment that when $d \neq 0$, delay until opt-out decreases approximately linearly with $|s - d|$, where s is the contribution rate initially chosen upon opt-out.

4.2 The Planner's Problem

We now describe the problem of a benign utilitarian social planner who sets the default savings rate d. The planner cannot observe each worker's optimal 401(k) savings rate s, although the planner knows the population distribution of s. For simplicity we assume that all workers have the same β and distribution of action costs, while s varies because of many unmodeled factors, including past and future bequests, past savings, spousal savings, historical (household) portfolio returns, and lifecycle family dynamics (e.g., children). We assume that s has a uniform density with support $[\underline{s}, \bar{s}]$ and renormalize the population so that the density is 1 at each point.

If the default is sufficiently far outside of $[\underline{s}, \bar{s}]$, then all workers are guaranteed to opt out of the default immediately. An active decision regime is mathematically equivalent to such a default within the model's framework, although in practice, active decisions are likely to be enforced mostly through social pressure rather than highly unattractive defaults.

Each individual's normative welfare measure ϕ is her expected utility loss discounted only by δ: $\phi_t = l_t + \delta l_{t+1} + \delta^2 l_{t+2} + \cdots$, where l_t is the expected loss she faces in period t. To motivate this long-run perspective, it is enough to assume that regulations established by the planner in period t take effect in period $t+1$. Then every worker at every point in time will want the planner to set the policy that minimizes ϕ. Henceforth, we set $\delta = 1$.

Without a deadline, if the worker's first cost draw is less than c^*, then her total realized loss is c; otherwise, she incurs an expected total loss of $L + \phi$ starting from the next period. So her expected total loss satisfies the equation

$$\phi = E(c \,|\, c < c^*)P(c < c^*) + (L + \phi)P(c > c^*). \tag{1}$$

We can write $\phi = \Phi(s - d)$, a function of $s - d$. The planner's problem is to choose d to

minimize

$$\int_{\underline{s}}^{\bar{s}} \Phi(s-d)\, ds. \tag{2}$$

This expression integrates the loss function for an individual with optimal savings rate s, which is $\Phi(s-d)$, over the population of all types $s \in [\underline{s}, \bar{s}]$.

4.3 Socially Optimal Default Policies

The following lemma states some properties of the loss function Φ and derives values of $s-d$ that will prove useful in defining candidate optimal default policies. Notice that $\Phi(s-d) = \Phi(-(s-d))$, since both arguments lead to the same L.

Lemma 1 *Let* $\Delta \equiv s-d$.

1. *If* $\beta = 1$, *then* $\Phi(\Delta)$ *is weakly increasing in* $|\Delta|$.

2. *If* $2 - \bar{c}/\underline{c} < \beta < 1$, *then (a) there exist* $0 < \Delta_m < \bar{\Delta}$ *such that* $\Phi(\Delta)$ *is increasing on* $(0, \Delta_m]$, *decreasing on* $[\Delta_m, \bar{\Delta}]$, *and constant at* $E(c) = (\underline{c} + \bar{c})/2$ *on* $[\bar{\Delta}, \infty]$, *and (b) there is at most one value* $\Delta_e \in (0, \bar{\Delta})$ *such that* $\Phi(\Delta_e) = E(c)$.

3. *If* $\beta \leq 2 - \bar{c}/\underline{c}$, *then there exists* $\bar{\Delta}$ *such that* $\Phi(\Delta)$ *is decreasing on* $(0, \bar{\Delta}]$ *and constant at* $E(c)$ *on* $[\bar{\Delta}, \infty]$.

Figure VII graphs total expected individual losses Φ as a function of Δ for various parameter sets. When $|\Delta|$ is sufficiently large, the worker will always immediately opt out of the default ($c^* = \bar{c}$) and therefore never incur any L flow losses. Thus the total loss is $E(c)$, which is independent of Δ. This is why the graphs are flat on the far left and right.

When $\beta = 1$ (the left graph in Figure VII), Φ is always weakly increasing as Δ moves away from zero; time-consistent workers are always weakly worse off with a larger flow loss L. But when $\beta < 1$ (the middle graph in Figure VII), there is an intermediate region in

which $\Phi(\Delta) > E(c)$. Workers in the "hump" of the loss function would be better off if L were much larger, since it would motivate them to overcome procrastination. If β is sufficiently low (the right graph in Figure VII), then *all* employees except those at the knife-edge value $\Delta = 0$ are made weakly better off by being forced to act immediately.

We can now characterize all possible optimal default policies.

Proposition 2 *If $\beta < 1$, then the optimal default is one of the following three types:*

1. *the* center default $d = (\underline{s} + \bar{s})/2$*;*
2. *an* offset default*, such that $\underline{s} - d = -\Delta_e$ and $\bar{s} - d \geq \bar{\Delta}$ (or its symmetric equivalent, $\bar{s} - d = \Delta_e$ and $\underline{s} - d \leq -\bar{\Delta}$);*
3. active decisions*, which correspond to any d with $\underline{s} - d \geq \bar{\Delta}$ or $\bar{s} - d \leq -\bar{\Delta}$.*

The possible optimal defaults when $\beta < 1$ correspond to the different panels of Figure VIII. The width of the shaded regions equals $\bar{s} - \underline{s}$, and their area equals the expected social cost associated with the respective default. The left panel shows the center default, where the default equals the mean of the s distribution. The middle panel shows the left offset default (there is also a symmetric right offset default). The offset default is placed so that workers with the lowest optimal savings rate, $\underline{s}$, opt out with some probability less than 1 in the first period, but procrastination causes their expected welfare loss to exactly equal the expected welfare loss if they were forced to opt out with certainty in the first period. The offset default also causes workers with the highest optimal savings rate, $\bar{s}$, to opt out with certainty in the first period because the default is far away from their optimum. If $\bar{s} - \underline{s}$ is not large enough to induce the two extremes of the population to behave in this manner, then the offset default is not a candidate for the optimal default. Finally, the right panel of Figure VIII shows an active decision regime, where the default is set so far outside of

the support of s that all individuals opt out of the default immediately and incur expected welfare losses of $E(c)$.

Figure IX shows how the optimal default depends on two parameters in our model: β (the level of time inconsistency) and $\bar{s} - \underline{s}$ (the heterogeneity of optimal savings rates). We discuss each region in turn.[23]

First consider the southeast region of Figure IX. In this region, employees have weak dynamic inconsistency problems and relatively homogeneous savings rates. The socially optimal solution here is a center default. This puts all employees in the middle of the graph of the function Φ, where the resulting losses are low.

As β falls, the humps in the graph of Φ grow, and eventually a center default puts so many workers in the humps that it is no longer optimal. The losses from procrastination become large relative to the option value of waiting for a low action cost, and employees are better off on average if they are forced to opt out of the default immediately. Thus, active decisions are optimal in this region.

When employees are relatively heterogeneous and β is high—so that the humps of Φ are not too large—the best solution is an offset default. Under a center default, employees would be in both humps, the valley, and possibly the plateaus. By using an offset default instead, the planner can beneficially move population mass from one of the humps to a plateau. The offset default is a compromise between the active decision and center default solutions. Some employees (but not all) are so far from their optimal savings rate that they are compelled to opt out of the default immediately, while others with optimal rates near the default are allowed to exploit the option value of waiting.

[23] When $\beta = 1$, the function Φ has no humps. When the range of savings rates is low, a center default is optimal. When savings rates are heterogeneous enough to cover the whole valley in the graph of Φ, then any default that is sufficiently far inside the interval $[\underline{s}, \bar{s}]$ to take full advantage of the valley will be optimal. See Proposition 14 in the online appendix.

The next proposition shows that the regions of Figure IX generically have the shape shown.

Proposition 3 *Fix κ, $\underline{c}$, and $\bar{c}$. Then there exist values $0 < \beta_{ac} < \beta_{oc} < 1$, and a function $w : (\beta_{ac}, 1] \to (0, \infty]$, with the following properties:*

1. *for $\beta \leq \beta_{ac}$, active decisions are always optimal;*
2. *for $\beta_{ac} < \beta < \beta_{oc}$, active decisions are optimal when $\bar{s} - \underline{s} > w(\beta)$ and a center default is optimal when $\bar{s} - \underline{s} < w(\beta)$;*
3. *for $\beta_{oc} < \beta < 1$, an offset default is optimal when $\bar{s} - \underline{s} > w(\beta)$ and a center default is optimal when $\bar{s} - \underline{s} < w(\beta)$;*
4. *w is increasing on $(\beta_{ac}, \beta_{oc}]$.*

5 Discussion

In this section, we interpret the model in light of real-world 401(k) enrollment mechanisms. We also discuss unmodeled factors that may drive opt-out delays or affect the welfare implications of enrollment mechanisms.

5.1 Interpreting the Model

We classify real-world 401(k) enrollment regimes into three types: standard enrollment, automatic enrollment, and active decisions. Under standard enrollment, employees have a default savings rate of zero and are given the option to raise this savings rate. Under automatic enrollment, employees have a default savings rate that is strictly positive and are given the option to change that savings rate (including opting out of the plan altogether).

Under active decisions, employees face no default and instead must affirmatively pick their savings rate.

Proposition 2 characterizes three types of optimal defaults, which correspond to these three regimes. The standard enrollment regime is an example of an offset default, since a 0% savings rate lies at one end of the optimal savings rate distribution.[24] The automatic enrollment regime, when implemented with a low contribution rate, is another example of an offset default.[25] In firms with higher default contribution rates, automatic enrollment is more like a center default.

It is perhaps surprising that the offset default is a candidate for an optimal default. Sunstein and Thaler (2003) express the intuition that a good default "minimizes the number of opt-outs." However, Proposition 2 shows that sometimes, the offset default—which is designed to encourage some rapid opt-outs—is welfare-maximizing. By using an offset default, the planner motivates some of the population to act quickly, avoiding large procrastination losses, while still letting others whose optimum is near the offset default exploit the option value of waiting to act.

When time inconsistency is weak (β is close to 1), a center default is optimal if employee savings preferences are relatively homogeneous. Beshears, Choi, Laibson, and Madrian (2008) discuss a number of techniques that a planner might use to estimate the distribution of s. For example, one could observe choices of workers who have had a sufficiently long time to move to their optimum. We believe that employees will have relatively homogeneous savings preferences when the workforce is demographically homogeneous (e.g., has a narrow

[24] The standard enrollment default is on the boundary of the action space, but it is conceptually similar to an offset default if a large number of households prefer not to save, while others have high optimal savings rates and would be motivated to act quickly by a zero default.

[25] The Profit Sharing/401(k) Council of America (2001) reports that before the Pension Protection Act of 2006, three-quarters of companies with automatic enrollment set their default contribution rate at 2% or 3% of pay, which is much lower than the 7% average 401(k) savings rate selected by employees when they make an affirmative choice (Holden and VanDerhei 2001).

range of ages) or when a generous employer match causes most employees to want to save at the match threshold.

As savings preferences become more heterogeneous, offset defaults such as standard enrollment and automatic enrollment with conservative defaults become more attractive. Standard enrollment and automatic enrollment with conservative defaults are also more attractive when a substantial fraction of employees have a low optimal savings rate in the 401(k). This may be the case if the company offers a generous defined benefit pension, if the company employs many low-wage workers who will have a high Social Security replacement rate, or if the company primarily employs young workers who would like to dissave at the present because they expect high future income growth.

If employees have a strong tendency to procrastinate (β is far below 1), then active decisions are optimal even when savings preferences have only a small amount of heterogeneity. The long-run stickiness of the default savings rate under automatic enrollment (Madrian and Shea 2001; Choi et al. 2004) supports the concern that many employees excessively delay opting out of defaults; it typically takes more than two years for the median employee to opt out of a 2 or 3% savings rate default. Active decisions eliminate the procrastination problem at the expense of losing the option value of waiting for a low-cost period to act.

The assumption that each individual's optimal savings rate is completely unobserved by the planner can be relaxed while preserving the results of Proposition 2 and 3. For example, consider the case in which the planner can calculate each employee's optimal savings rate s, but the employer is constrained to implement the same default policy for the entire company. Our model's results also continue to hold if employee-specific defaults *are* feasible but the planner observes s with noise. Then the model's optimal savings rate heterogeneity can be reinterpreted as the heterogeneity within a subpopulation that will share a common default, and the optimal default generated by the model is the optimal default for that particular

subpopulation.

5.2 Financial Illiteracy and Defaults

There is a growing body of evidence that planners make better asset allocation choices than ordinary households (Benartzi and Thaler 2001; Cronqvist and Thaler 2004; Choi, Laibson, and Madrian 2008). One manifestation of households' ignorance is their interpretation of and reliance upon asset allocation defaults as implicit advice from the company, even when the default is not intended to serve such a purpose. Madrian and Shea (2001) find that employees not subject to automatic enrollment shift their 401(k) asset allocations towards the default fund after their company automatically enrolls *other* employees.

Our model assumes workers have better information about their optimum than the planner and that the planner's choice of a default does not affect workers' perception of what that optimum is. These assumptions are violated in the domain of asset allocation. Therefore, the welfare results of our model do not necessarily apply to asset allocation choices.

On the other hand, workers are probably better-informed about their optimal savings rate than planners. When asked what their ideal retirement savings rate was, workers in the Choi et al. (2002) survey gave an answer of 14% on average, which is close to what many financial experts would recommend. Two-thirds of surveyed households recognized that they were saving too little, suggesting that widespread undersaving is not primarily driven by ignorance about the need to save. Moreover, there was dispersion in reported optimal savings rates, as should be expected given idiosyncratic savings needs due to expected income growth, existing financial assets, years until retirement, children in the household, etc. Consistent with greater household knowledge about optimal savings rates, in the company studied by Madrian and Shea (2001), workers do not appear to take the employer-chosen contribution rate default as advice under automatic enrollment; workers not subject to automatic enroll-

ment do not disproportionately shift their 401(k) contribution rates to the default applied to automatically enrolled employees. Similarly, Choi et al. (2004) find that automatically enrolled workers are more likely to opt out of contribution rate defaults than asset allocation defaults.

In summary, our model is better applied to savings rate choices rather than asset allocation choices.[26] Well-chosen defaults are likely to be superior to active decisions in the asset allocation domain. Active decisions can easily be implemented for savings rate choices while maintaining an asset allocation default.[27]

5.3 Absent-Mindedness and Opt-Out Delays

In our model, employees act as if each period, they were consciously comparing the cost of opting out to the delayed benefit from opting out. In reality, people are likely to (endogenously) go through long stretches of time when they do not think about their 401(k). If such absent-mindedness is a significant driver of opt-out delays, merely reminding employees about their 401(k) may be enough to motivate significant action.

To assess the efficacy of reminders, we conduct new analysis on survey data collected in the field experiment run by Choi, Laibson, and Madrian (2007). Surveys were mailed to 1,503 employees at one company who were not enrolled in their 401(k). Among other things, the survey asked respondents how much they were actually saving for retirement, how much they should ideally be saving, and when (if ever) they planned to enroll in their 401(k).[28]

[26] A concern specific to savings rate choices is that present-biased agents will generally want to undersave. But to the extent that an active 401(k) contribution rate choice is a commitment to save in the *future*, starting with the next paycheck, present-biased agents at the point of action will choose the optimal (from the long-run self's perspective) contribution rate.

[27] For example, in an online active decision implementation, employees would be required to actively enter a contribution rate. They would be told that any contributions will be invested in a default asset allocation unless they click a link to change that asset allocation.

[28] This sample includes workers both younger and older than $59\frac{1}{2}$. The 1,503 employees we analyze here do not include those who were given a treatment survey designed to educate them about the arbitrage opportunity they were leaving unexploited by not contributing up to the employer match threshold. See

Most of the 1,380 survey recipients who did not respond to our survey probably never read it and hence were not reminded of their 401(k). We see in the 401(k) administrative data that 3.7% of these non-responders enrolled in the plan during the four months following the survey. Among the 123 survey recipients who did respond to our survey, only 3.3% enrolled in the subsequent four months. Both groups of unenrolled workers showed little propensity to enroll despite being mailed our implicit reminder. Even more surprising, the group that was definitely reminded (survey responders) had a slightly lower enrollment rate than the group that probably was not reminded (non-responders). Thus, it seems that merely causing people to think about their 401(k) has little impact on subsequent enrollment behavior.

5.4 Manipulating Opt-Out Costs

We have modeled the distribution of opt-out costs c as a fixed, exogenous parameter. In practice, the c distribution can be affected by the planner. For example, the employer can make the process of opting out more or less inconvenient. Employers can also decrease opt-out costs by making 401(k) decisions cognitively simpler.

Choi, Laibson, and Madrian (forthcoming) and Beshears et al. (2006) study an intervention called Quick Enrollment, which works on both the convenience and cognition margins to reduce opt-out costs. In one implementation, employees who were not enrolled in their 401(k) plan received a postcard with a pre-selected asset allocation and a pre-selected contribution rate. If an employee wished to enroll with those elections, the employee simply signed the card and dropped it back in the mail (its postage was pre-paid). In another implementation, employees were given the opportunity to enroll in the 401(k) via a special Web interface. Employees chose their own contribution rate, but could elect to invest contributions to a pre-selected asset allocation. In both implementations, employees remained free to enroll

Choi, Laibson, and Madrian (2007) for further details on this educational treatment.

at elections other than the pre-selected ones by using standard phone or Internet channels. Employees who used Quick Enrollment could also change their elections afterwards using the standard channels.

Quick Enrollment increases participation rates by 10 to 25 percentage points relative to a regime with a non-enrollment default and no pre-selected enrollment options. It does so by reducing the cost of opting into the plan in two ways. First, the card or special Web interface may just be a more convenient way to enroll. Second, Quick Enrollment simplifies a complicated multi-dimensional problem. In the first implementation, employees could focus on making a binary choice between the status quo and the pre-selected investment plan, rather than considering the entire menu of potential 401(k) elections. In the second implementation, employees could make the contribution rate decision without making an asset allocation decision. Ordinarily, a newly enrolling employee must simultaneously make both decisions. Decoupling the decisions reduces the cognitive cost of enrollment, since (as discussed in Section 5.2) households tend to have more knowledge about optimal savings rates than optimal asset allocations.

Quick Enrollment can be combined with any default or active decision regime. Quick Enrollment can also be extended to offer several options; for example, the Quick Enrollment card could ask employees to select either the Conservative, Moderate, or Aggressive asset allocation. As the number of Quick Enrollment menu options increases, complexity also rises, causing action costs to approach those of a regime without Quick Enrollment. Quick Enrollment will be especially useful when decision-making costs are high and a small number of choices (corresponding to the options on the Quick Enrollment menu) are close to the investment optima of a large fraction of the population.

6 Conclusion

This paper identifies and analyzes the active decision alternative to default-based 401(k) enrollment processes. The active decision approach forces employees to explicitly choose between the options of enrollment and non-enrollment in the 401(k) plan without advantaging either of these outcomes. We find that the fraction of employees who enroll in the 401(k) three months after hire is 28 percentage points higher under an active decision regime than under a standard opt-in enrollment regime. The active decision regime also raises average saving rates and accumulated 401(k) balances. The distribution of employees' savings rates immediately after hire under active decisions is similar to the distribution observed thirty months after hire under standard enrollment.

These results suggest that low 401(k) participation rates under standard enrollment regimes do not entirely reflect a desire to save nothing in the 401(k). When forced to actively choose early in their tenure, most employees choose to contribute. The large participation increase under active decisions is also evidence that much of the participation increases under automatic enrollment are not due to employers "tricking" employees into joining the plan.

We present a general model of procrastination which can be used to describe the employee's 401(k) enrollment choice. Under this framework, we describe conditions under which the optimal enrollment regime is automatic enrollment (i.e., default enrollment), standard opt-in enrollment (i.e., default non-enrollment), or active decisions (i.e., no default and compulsory choice). The active decision regime is socially optimal when consumers have relatively heterogeneous savings preferences and a strong tendency to procrastinate.

In the current technological environment, an active decision enrollment regime should not take the form of the paper-and-pencil system that we studied. Instead, active decision systems could be designed to take advantage of the efficiencies available with electronic enrollment. For example, a firm could require new employees to visit a Web site where

they would actively elect to enroll in or opt out of the 401(k) plan, perhaps in conjunction with electing other benefits or providing other information relevant to the company.[29] Firms could also compel non-participating employees to make an active decision during each annual open enrollment period for other benefits (e.g. health insurance). This would ensure that non-participating employees rethink their non-participation in the 401(k) at least once a year.

The active decision approach to increasing 401(k) participation has some attractive features relative to other savings schemes. Active decision is a relatively low cost way to boost 401(k) participation (e.g. active decisions are much less expensive than matching and far more effective in raising employee contributions). Requiring individuals to make an active decision represents a weaker alternative to the standard paternalism implicit in specifying a default. Active decision interventions are designed principally to force a decision-maker to think about a problem. This is still a type of paternalism, but it does not presuppose an answer to the problem.

We should note that we are not opposed to other savings interventions, including financial education, employer matching, or automatic enrollment. We view all of these, along with active decisions, as complementary approaches to fostering increased and higher-quality 401(k) participation.

Active decision interventions will be useful in many other situations where consumer heterogeneity implies that one choice isn't ideal for everyone (e.g., the selection of a health plan or automobile insurance[30]) and firms or governments feel uncomfortable implementing employee-specific defaults (e.g., if such employee-specific defaults are viewed as "advice" with

[29] Workers without access to computers could submit paper forms.

[30] The active decision approach to purchasing automobile insurance is widely used. Drivers cannot, in general, register their cars without obtaining insurance. But the government does not specify a default insurance contract for drivers; rather, it requires drivers to obtain their own insurance—to make an active decision. The model in the paper suggests that there is a good justification for this approach: there is likely to be substantial heterogeneity in individual preferences over insurance policy types and companies.

fiduciary consequences). In contrast, defaults will have a natural role to play in cases where consumers are relatively homogenous and household decision-makers have limited expertise (e.g., portfolio allocation). Future research should explore active decision experiments in other decision domains and compare the relative efficacy of active decision and default-based systems, as well as hybrid systems which will probably turn out to be the most useful approach of all.

Massachusetts Institute of Technology

Yale University and NBER

Harvard University and NBER

Harvard University and NBER

Yale University and NBER

7 References

Abadie, Alberto, and Sebastien Gay, "The Impact of Presumed Consent Legislation on Cadaveric Organ Donation: A Cross Country Study," *Journal of Health Economics*, 25 (2004), 599–620.

Benartzi, Shlomo, and Richard H. Thaler, "Naive Diversification Strategies in Defined Contribution Saving Plans," *American Economic Review*, 91 (2001), 79–98.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, "Simplification and Saving," NBER Working Paper No. 12659, 2006.

——, "How Are Preferences Revealed?," *Journal of Public Economics*, 92 (2008), 1787–1794.

Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and Brian Weller. "Public Policy and Saving for Retirement: The 'Autosave' Features of the Pension Protection Act of 2006," in *Better Living Through Economics: How Economic Research Improves Our Lives*, John J. Siegfried, ed. (Cambridge, MA: Harvard University Press, forthcoming).

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Optimal Defaults and Active Decisions," NBER Working Paper No. 11074, 2007.

Choi, James J., David Laibson, and Brigitte C. Madrian, "Reducing the Complexity Costs of 401(k) Participation Through Quick Enrollment^{TM}," in *Developments in the Economics of Aging*, David A. Wise, ed. (Chicago: University of Chicago Press, forthcoming).

——, "$100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans," NBER Working Paper No. 11554, 2007.

——, "Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds." NBER Working Paper 12261, 2008.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Defined Contribution Pensions: Plan Rules, Participant Decisions, and the Path of Least Resistance," in *Tax Policy and the Economy* 16, James M. Poterba, ed. (Cambridge, MA: MIT Press, 2002).

——, "Optimal Defaults," *American Economic Review*, 93 (2003), 180–185.

——, "For Better or For Worse: Default Effects and 401(k) Savings Behavior," in *Perspectives in the Economics of Aging*, David A. Wise, ed. (Chicago, IL: University of Chicago Press, 2004).

Cronqvist, Henrik, and Richard H. Thaler, "Design Choices in Privatized Social-Security Systems: Learning from the Swedish Experience," *American Economic Review*, 94 (2004), 424–428.

Even, William E., and David MacPherson, "The Effects of Employer Matching in 401(k) Plans," *Industrial Relations*, 44 (2005), 525–549.

Harris Interactive, "Retirement Made Simpler," downloaded from http://www.retirementmadesimpler.org/Library/FINAL%20RMS%20Topline%20Report%2011-5-07.pdf on December 28, 2007.

Hewitt Associates, *Trends and Experience in 401(k) Plans: 2007* (Lincolnshire, IL: Hewitt Associates, 2007).

Holden, Sarah, and Jack VanDerhei, "Contribution Behavior of 401(k) Plan Participants," *Investment Company Institute Perspective*, 7 (2001).

John Hancock, "Eighth Defined Contribution Plan Survey: Insight Into Participant Knowledge & Behavior." (Boston, MA: John Hancock Financial Services, 2002).

Johnson, Eric J., and Daniel G. Goldstein, "Do Defaults Save Lives?," *Science*, 302 (2003), 1338–1339.

Johnson, Eric J., John Hershey, Jacqueline Meszaros, and Howard Kunreuther, "Framing, Probability Distortions, and Insurance Decisions," *Journal of Risk and Uncertainty*, 7 (1993), 35–53.

Johnson, Eric J., Steven Bellman, and Gerald L. Lohse, "Defaults, Framing and Privacy: Why Opting In-Opting Out," *Marketing Letters*, 13 (2002), 5–15.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler, "The Endowment Effect, Loss Aversion, and Status Quo Bias," *Journal of Economic Perspectives*, 5 (1991), 193–206.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112 (1997), 443–477.

Madrian, Brigitte C., and Dennis Shea, "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics*, 116 (2001), 1149–1187.

O'Donoghue, Ted, and Matthew Rabin, "Procrastination in Preparing for Retirement," in *Behavioral Dimensions of Retirement Economics*, Henry J. Aaron, ed. (Washington, D.C.: Brookings Institution Press; New York, NY: Russell Sage Foundation, 1999a).

O'Donoghue, Ted, and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 89 (1999b), 103–124.

Park, C. Whan, Sung Y. Jun, and Deborah J. MacInnis, "Choosing What I Want Versus Rejecting What I Do Not Want: An Application of Decision Framing to Product Option Choice Decisions," *Journal of Marketing Research*, 37 (2000), 187–202.

Phelps, Edmund S., and Robert A. Pollak, "On Second-Best National Savings and Game-Equilibrium Growth," *Review of Economic Studies*, 35 (1968), 185–199.

Profit Sharing/401(k) Council of America, "Automatic Enrollment 2001: A Study of Automatic Enrollment Practices in 401(k) Plans," downloaded from http://www.psca.org/data/autoenroll2001.asp on April 19, 2001.

Samuelson, W., and R. J. Zeckhauser, "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1 (1988), 7–59.

Skinner, Jonathan, "Are You Sure You're Saving Enough for Retirement?," *Journal of Economic Perspectives*, 21 (2007), 59–80.

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

Table I. 401(k) plan features by effective date

	Effective January 1, 1997	Effective November 23, 1997
Eligibility		
Eligible employees	U.S. employees, age 18+	U.S. employees, age 18+
First eligible	Full-time employees eligible upon hire; part-time employees must accrue 1,000 hours in one year	Full-time employees eligible upon hire; part-time employees must accrue 1,000 hours in one year
Employer match eligible	Immediately upon plan eligibility	Immediately upon plan eligibility
Enrollment	First thirty days of employment or January 1 of succeeding calendar years	Daily
Contributions		
Employee contributions [a]	Up to 17% of compensation	Up to 17% of compensation
Guaranteed employer match	50% of employee contribution	50% of employee contribution
Additional possible employer match	Up to 100% (50% for bonus-eligible employees); rate depended on company profitability	Up to 100% (50% for bonus-eligible employees); rate depended on company profitability[b]
Employer match restrictions	Match on the lesser of before-tax employee contribution or 5% of compensation; match invested in employer stock	Match on the lesser of before-tax employee contribution or 5% of compensation; match invested in employer stock
Employer match vesting	Immediate	Immediate
Other		
Loans	Not available	Available; two maximum
Hardship withdrawals	Available	Available
Investment choices	Four options; employer stock also available, but only for after-tax contributions and employer match	Six options + employer stock (available for before- and after-tax contributions)

[a]Total employee contributions within each year were capped by federal law at $9,500 (1997), $10,000 (1998-99), and $10,500 (2000-01).

[b]Actual discretionary match rates were 20% (1995), 20% (1996), 100% (1997), 100% (1998), 27% (1999), 33% (2000), 0% (2001).

Table II. Comparison of worker characteristics

	Study company			
	Active decision cohort on 12/31/98	Standard enroll. cohort on 12/31/99	All workers on 12/31/99	U.S. workforce (3/98 CPS)
Average age (years)	34.7	34.1	40.8	38.8
Gender				
Male	47.6%	42.0%	44.7%	53.1%
Female	52.4%	58.0%	55.3%	46.9%
Marital Status				
Single	41.4%	49.3%	32.2%	39.0%
Married	56.0%	49.7%	66.8%	61.0%
Compensation				
Avg. monthly base pay	$3,043	$2,869	$4,367	--
Median monthly base pay	$2,666	$2,513	$3,664	--
Avg. annual income[a]	$35,381	$33,197	$50,414	$32,414
Median annual income[a]	$31,013	$29,239	$40,965	$24,108
Geography				
East	13.2%	11.1%	15.0%	18.9%
Midwest	34.3%	37.6%	32.2%	24.1%
South	37.7%	38.9%	37.7%	34.7%
West	14.7%	12.3%	15.0%	22.4%
Number of Employees	2,231	2,349	46,944	--

The samples in the first three columns are taken from individuals employed at the study company as of the dates indicated in the column title. The sample in the last column is all individuals (weighted) in the March 1998 Current Population Survey who worked in the previous year. Compensation is in 1998 dollars. Figures may not add up to 100% because of missing data and employees located in Puerto Rico.

[a]The annual income measure that is reported to us for the study company is the employee's annual taxable (W2) income. Annual income for the U.S. workforce calculated from the CPS is total annual labor earnings in the previous calendar year, some of which may be non-taxable.

Table III. Tobit regression of contribution rates under two 401(k) enrollment regimes	
Intercept	2.915**
	(0.191)
Female	0.545
	(0.403)
Married	1.335**
	(0.383)
Log(*Base pay*)	4.898**
	(0.666)
$0 \leq Age < 30$	-6.829*
	(2.681)
$30 \leq Age < 40$	-5.144
	(2.683)
$40 \leq Age < 50$	-5.133
	(2.703)
$50 \leq Age < 60$	-3.182
	(2.753)
Active decision cohort	0.086
	(0.247)
Active decision cohort × *Female*	-1.989**
	(0.547)
Active decision cohort × *Married*	-0.528
	(0.503)
Active decision cohort × Log(*Base pay*)	-1.930
	(1.053)
Active decision cohort × $(0 \leq Age < 30)$	-0.584
	(3.553)
Active decision cohort × $(30 \leq Age < 40)$	-0.577
	(3.552)
Active decision cohort × $(40 \leq Age < 50)$	0.004
	(3.578)
Active decision cohort × $(50 \leq Age < 60)$	-0.103
	(3.659)
ln $(\sigma_{AD}/\sigma_{SE})$	-0.137**
	(0.042)
N	3,488

If the employee is in the active decision cohort, the dependent variable is the estimated 401(k) contribution rate (in percentage points) three months after hire; if the employee is in the standard enrollment cohort, the dependent variable is the contribution rate thirty months after hire. Independent variables are log of base pay, a dummy for being in the active decision cohort, and gender, marital status, and age range dummies, calculated as of the contribution rate date. Demographic variables are de-meaned. Both cohorts are restricted to employees who remain in the data for at least thirty months. The tobit regression assumes that errors are normal and homoskedastic within each cohort but possibly heteroskedastic across cohorts. Robust standard errors are reported in parentheses under the point estimates.

*Significant the 5% level **Significant at the 1% level

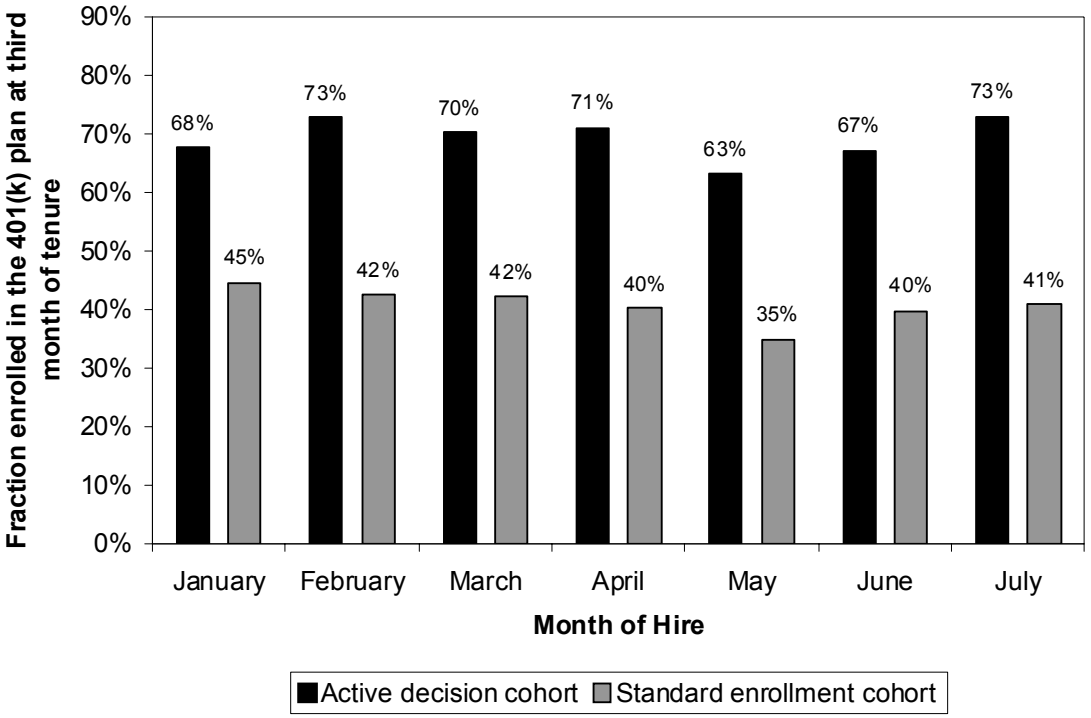

Figure I

Fraction of Employees Enrolled in the 401(k), by Hire Month

The fraction displayed is as of the third month of tenure at the company. The active decision cohort was hired between January and July 1997. The standard enrollment cohort was hired between January and July 1998.

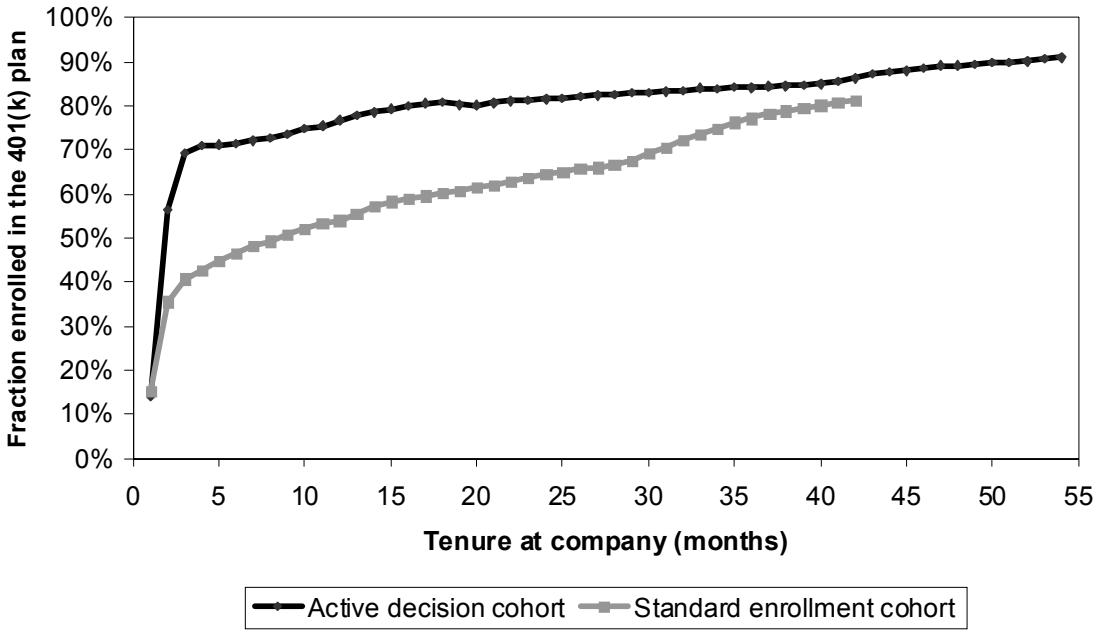

Figure II

Fraction of Employees Enrolled in the 401(k) Plan, by Tenure at Company

An employee is counted enrolled in the 401(k) even if he or she has stopped contributing to the plan. The series are not monotonically rising because they are constructed from multiple cross-sections, so the samples are not fixed over time.

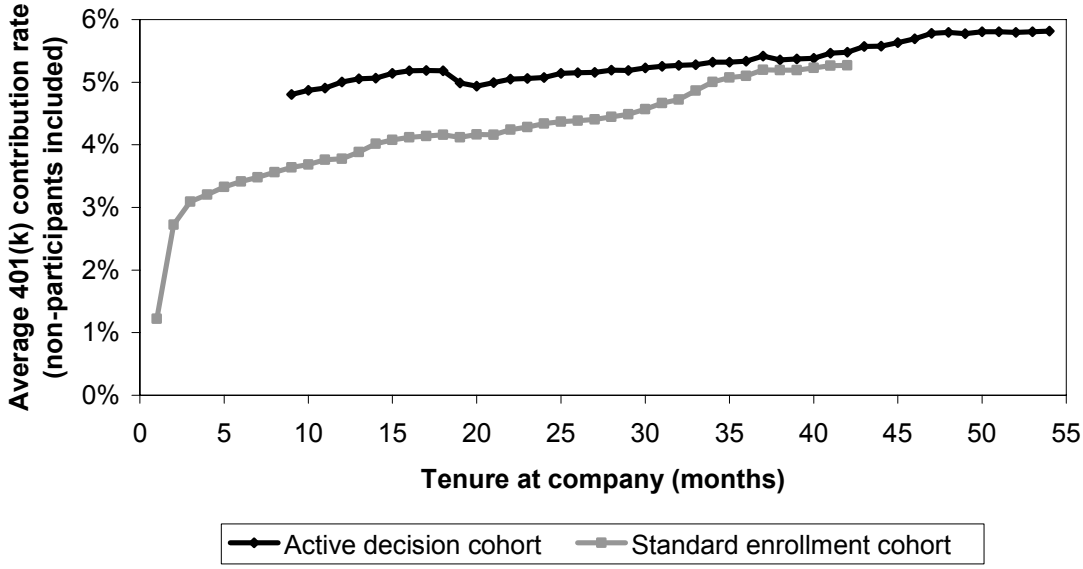

Figure III

Average 401(k) Contribution Rate, by Tenure at Company

At each point, the averages include employees not currently contributing to the 401(k) plan; their contribution rate is zero. The active decision cohort's contribution rate data are not available prior to month nine.

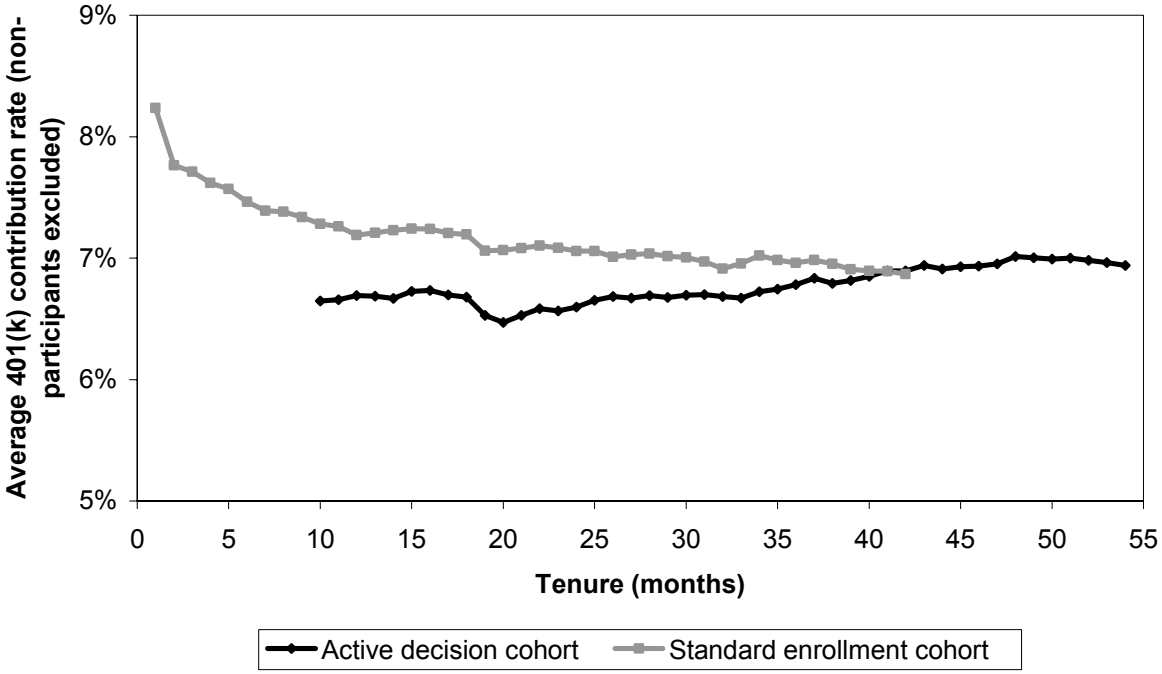

Figure IV

Average 401(k) Contribution Rate Conditional on Participating, by Tenure at Company

At each point, the averages exclude employees not currently contributing a positive amount to the 401(k) plan. The active decision cohort's contribution rate data are not available prior to month nine.

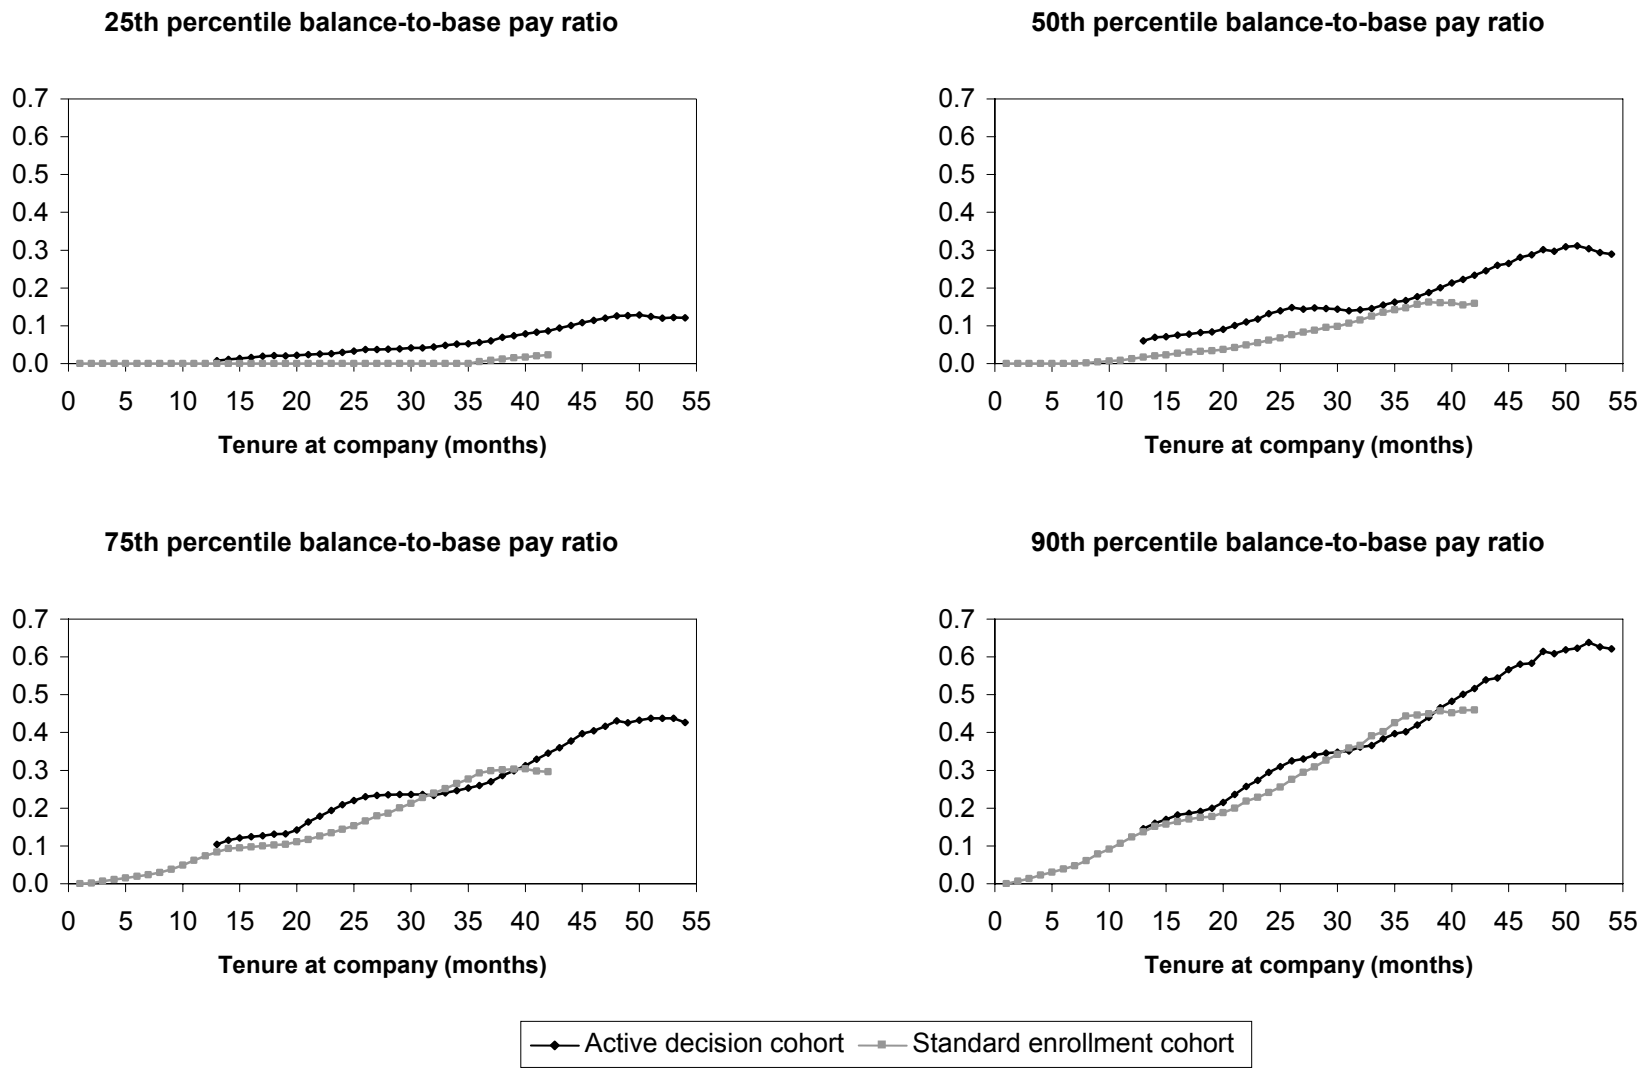

Figure V

401(k) Balance-to-Base Pay Ratios at Different Balance-to-Base Pay Percentiles

Balances exclude outstanding loan principal and any money rolled into the account from a former employer. The percentile breakpoints are calculated separately for each cohort at each point in time. The active decision series starts in month thirteen because our salary data start in January 1998.

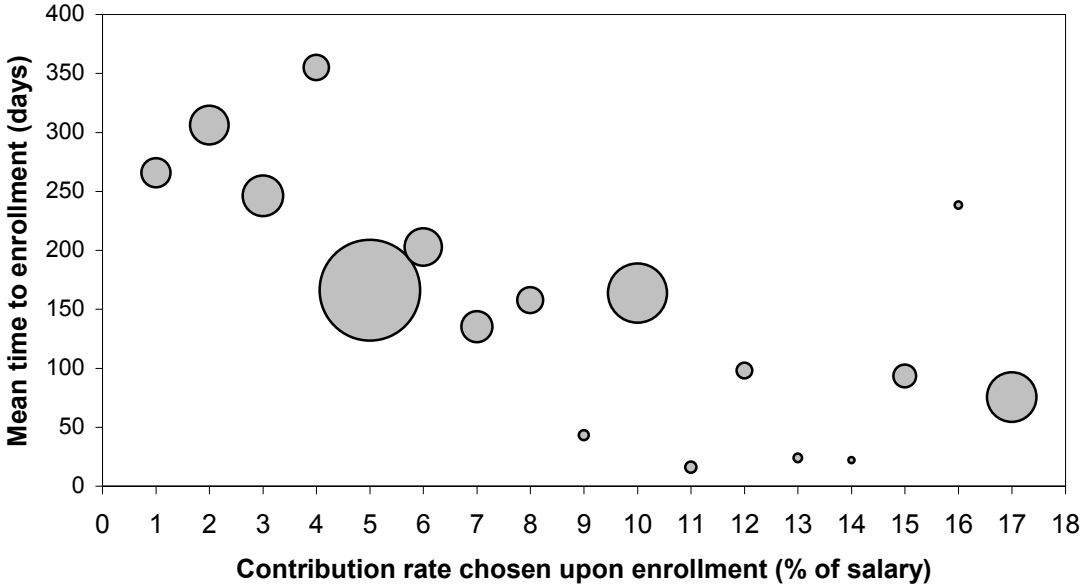

Figure VI

Mean Time Between Hire Date and Enrollment by

Contribution Rate Chosen Upon Enrollment in the Standard Enrollment Cohort.

The area of each bubble is proportional to the number of employees it represents. The sample consists of all employees in the standard enrollment cohort who worked at the company for at least thirty months and enrolled within thirty months of hire.

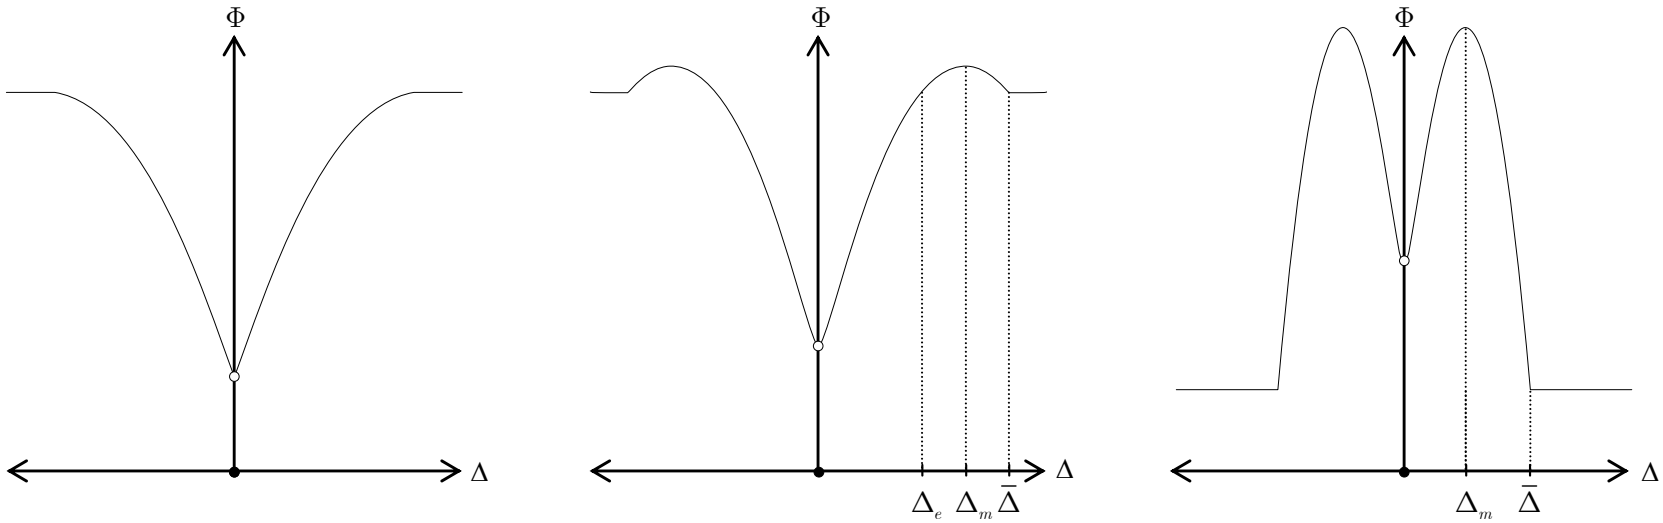

$\beta = 1$ $\beta = 0.75$ $\beta = 0.1$

Figure VII

Employee's Total Expected Loss Φ as a Function of the Distance Between the Default and the Employee's Optimum

The parameters specific to each panel, the quasi-hyperbolic discount factor β, is shown beneath each graph. In all three panels, the opt-out cost is uniformly distributed between 0.25 and 1.75, and the loss function scaling factor $\kappa = 100$. The left and center panels have the same y-axis scale, but the right panel has a different y-axis scale.

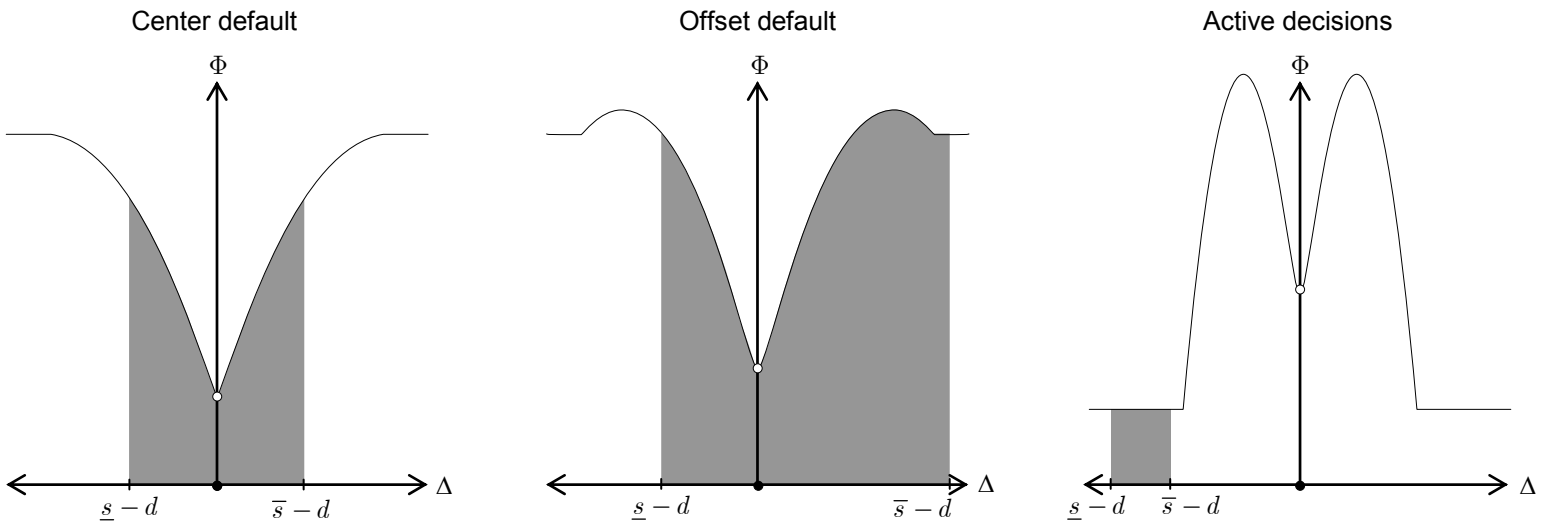

Figure VIII

Possible Optimal Default Regimes

The panels illustrate parameter values that support the three classes of optimal defaults: the center default, the offset default, and active decisions. The shaded area in each panel represents the social welfare losses generated by the corresponding default regime. The parameters specific to each panel, the quasi-hyperbolic discount factor β and the range of optimal savings rates $\overline{s} - \underline{s}$, appear below each figure. In all three panels, the opt-out cost is uniformly distributed between 0.25 and 1.75, and the loss function scaling factor $\kappa = 100$. The left and center panels have the same y-axis scale, but the right panel has a different y-axis scale.

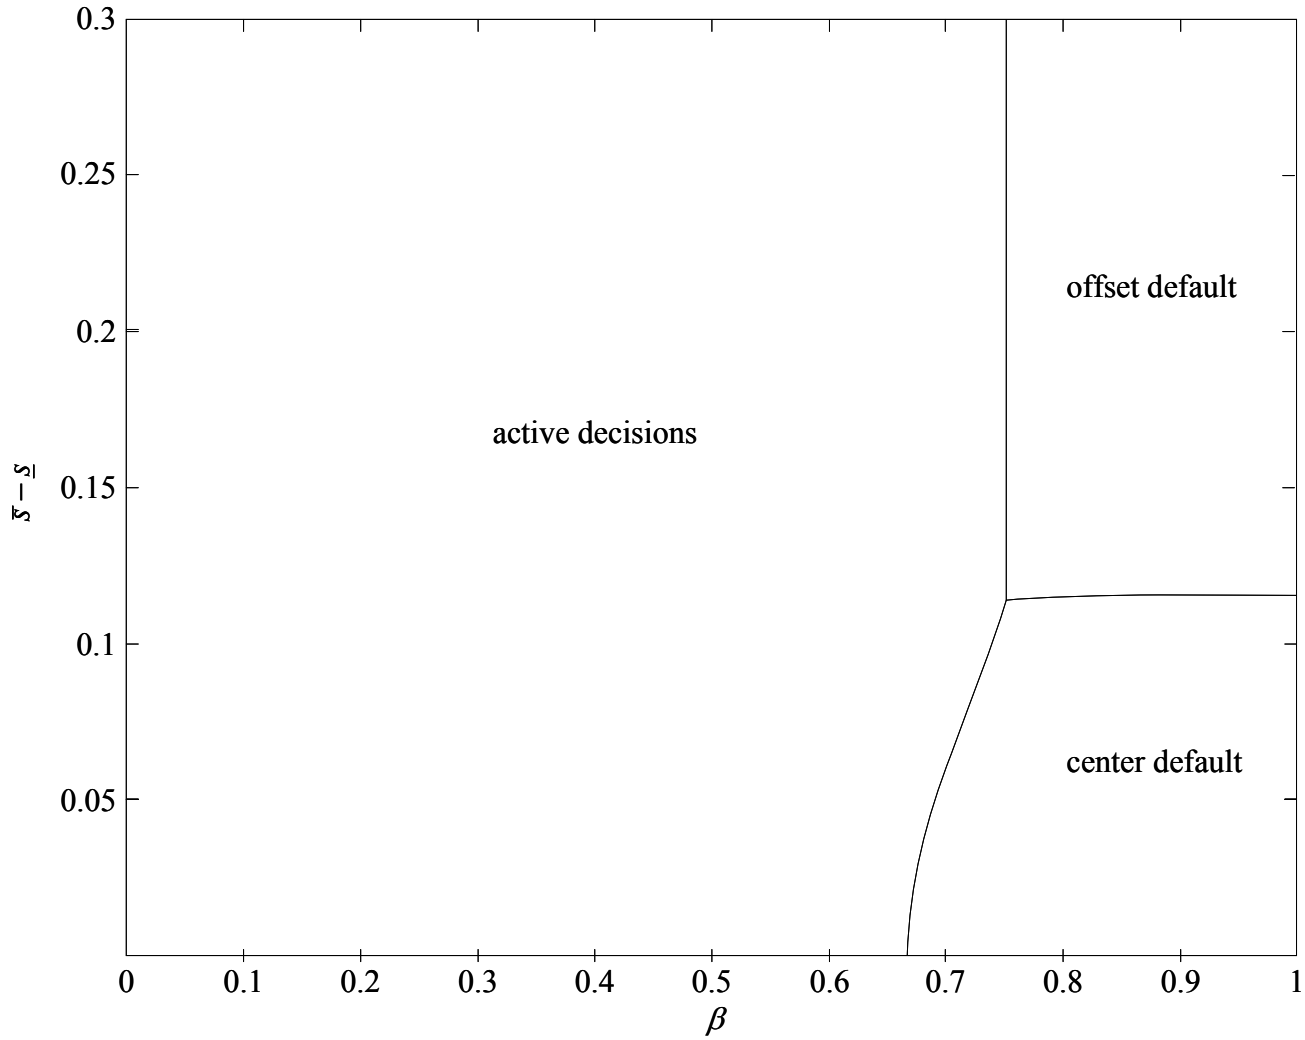

Figure IX

Characterization of Optimal Default Regimes

This figure shows the boundaries of the optimal default regimes as a function of the quasi-hyperbolic discount factor β and the range of optimal savings rates $\overline{s} - \underline{s}$ when the opt-out cost is uniformly distributed between 2/3 and 4/3, and the loss function scaling factor $\kappa = 100$.