

NBER WORKING PAPER SERIES

PASSIVE DECISIONS AND POTENT DEFAULTS

James J. Choi
David Laibson
Brigitte C. Madrian
Andrew Metrick

Working Paper 9917
<http://www.nber.org/papers/w9917>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2003

This paper is a longer version of a paper published in the *American Economic Review*, May 2003 (Papers and Proceedings). We thank Hewitt Associates for their help in providing the data. We are particularly grateful to Lori Lucas, Jim McGhee, and Scott Peterson, three of our many contacts at Hewitt. We are grateful for the comments of Robert Barro, Robert Hall, Mark Iwry, Antonio Rangel, Richard Thaler, and seminar participants at Harvard, MIT, NBER, University of Chicago, University of Minnesota, University of Southern California, and Wharton. Numerous research assistants have contributed to this work. We are particularly indebted to Nelson Uhan. We acknowledge financial support from the National Institute of Aging (grant R01AG021650). Choi acknowledges support from an NSF Graduate Research Fellowship. Laibson acknowledges financial support from the National Science Foundation (grant 0099025). The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

©2003 by James J. Choi, David Laibson, Brigitte Madrian, and Andrew Metrick. 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.

Passive Decisions and Potent Defaults

James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick

NBER Working Paper No. 9917

August 2003

JEL No. D1

ABSTRACT

Default options have an enormous impact on household "choices." Defaults matter because opting out of a default is costly and these costs change over time, generating an option value of waiting. In addition, people have a tendency to procrastinate. We develop a theory of optimal defaults based on these considerations. We find that it is sometimes optimal to set extreme defaults, which are far away from the mean optimal savings rate. A default that is far away from a consumer's optimal savings rate may make that consumer better off since such a "bad" default will lead procrastinating consumers to more quickly opt out of the default. We calibrate our model and use it to calculate optimal defaults for employees at four different companies. Our work suggests that optimal defaults are likely to be at one of three savings rates: the minimum savings rate (i.e., 0%), the match threshold (typically 5% or 6%), or the maximal savings rate.

James Choi
Department of Economics
Harvard University
Cambridge, MA 02138
james_choi@post.harvard.edu

David Laibson
Department of Economics
Harvard University
Cambridge, MA 02138
and NBER
dlaibson@harvard.edu

Brigitte Madrian
Wharton School
University of Pennsylvania
Philadelphia, PA 19104
and NBER
bmadrian@wharton.upenn.edu

Andrew Metrick
Wharton School
University of Pennsylvania
Philadelphia, PA 19104
and NBER
andrew.metrick@wharton.upenn.edu

Default options have an enormous impact on household “choices.” Such effects have now been extensively documented in the literature on 401(k) plans.¹ Defaults have been shown to affect participation, savings rates, rollovers, and asset allocation. For example, Choi et al. (2003b) study three firms that use automatic enrollment. When employees at these firms are automatically enrolled in their 401(k) plan, only a tiny fraction opt out, producing participation rates exceeding 85% regardless of tenure. But when employees at these firms were not automatically enrolled, participation rates were significantly lower, ranging from 26% – 43% after six months of tenure, and from 57% – 69% after three years of tenure.

Defaults matter for three key reasons that we model in this paper. First, acts of commission — e.g., opting out of a default — are costly. Second, these costs change randomly over time and therefore generate an option value of waiting to change a default. Decision makers would like to wait for a low cost period (e.g., a free weekend) to make a change. Third, people have a tendency to procrastinate. Even if they want to make a change, they have a tendency to delay that change longer than they should.

Because of these effects, the choice of a particular default can have a significant effect on consumer welfare. However, it is not always obvious how to select a socially optimal default.

If *all* employees would like to be saving at a rate of exactly 5% in their 401(k) plan, then the employees’ welfare will be maximized if the employer sets a 5% default. But

¹Madrian and Shea (2001), Choi et al (2002a, 2003b).

the calculation of an optimal default is not as straightforward if different employees have different optimal savings rates. For example, what is the optimal default savings rate if employees have optimal savings rates that are distributed uniformly with a mean of 5%?

In this paper, we develop a theory of optimal defaults which implies that the obvious answer to the previous question — 5% — is not necessarily the right answer. In a world of heterogeneous agents, it may sometimes be optimal to set *extreme* defaults which are far away from the mean optimal savings rate. This effect arises for two reasons. First, a default that is far away from a consumer’s optimal savings rate may make that consumer *better off* than a default that is closer to the consumer’s optimal savings rate. Intuitively, if an agent suffers from a procrastination problem, then a “bad” default — i.e., one that is far from the consumer’s optimal savings rate — will be more motivating than a better default. Hence, sometimes bad defaults make people better off than better but imperfect defaults. Second, our theory implies that optimal defaults are highly sensitive to the actual distribution of optimal savings rates. In particular, optimal defaults are often associated with the modal optimal savings rate and not the mean optimal savings rate. Since these modes are sometimes extreme (e.g., minimum or maximum contribution rates), optimal defaults will sometimes be extreme as well.

At the end of our paper we calibrate our model and use it to calculate optimal defaults for employees at four different companies. For two of our companies, the optimal default is close to the mean optimal savings rate, whereas for the other two companies

the optimal defaults are extreme: 0% and 15% respectively. Our work suggests that optimal defaults are likely to be at one of three savings rates: the minimum savings rate (0%), the employer match threshold (typically 5% or 6%), or the maximal savings rate (around 15% in our sample of companies from the late 1990's).²

1. A MODEL OF SAVINGS CHOICES

We adapt the model of Choi et al. (2002b) to describe the 401(k) enrollment decisions of employees that have been newly hired at a firm. However, the model is general enough to describe any problem in which an actor decides when to move from a default state s_D to an optimal state s^* .

We assume that each employee at a firm has a fixed optimal savings rate (i.e., optimal state) s^* , with density function f characterizing the distribution of these optimal savings rates for the population of employees in the firm. When new employees join the firm, the employees are automatically enrolled at a default savings rate of s_D , which is a choice variable for the firm. In this paper, we consider the case in which this default can only take values in the support of f .³ We assume that the firm uses a single default savings rate for all of its employees either because the firm does not observe an employee's true type, s^* , or because of legal/practical costs of implementing employee-specific defaults.⁴

²More recently, regulatory changes under the Economic Growth and Tax Relief Reconciliation Act of 2001 (EGTRRA) have led many companies to raise their maximum savings rates well above the historical norm of 15%.

³See Choi et al. (2002b) for a generalization.

⁴Such employee-specific defaults are a natural extension of our current framework and merit theoretical and practical evaluation.

Employees remain at the default election s_D unless they opt out of the default by incurring a cost c . This opt-out cost is drawn each period and takes the value 1 with probability $\mu < 1$ and value 0 with probability $1 - \mu$. The value of the cost is known when the agent decides on her action. We suppress individual and time subscripts to simplify notation.

When the agent opts out, she sets her savings rate equal to her optimal savings rate s^* , which we assume the agent knows with certainty.⁵ Until that action takes place, the agent suffers a flow loss of $L = L(s_D, s^*) \geq 0$, where the first argument of L is the current savings rate and the second argument of L is the optimal savings rate. After the action occurs, the agent suffers a flow loss of $0 = L(s^*, s^*)$.

Finally, we assume that agents are *naive* hyperbolic discounters, with discount function $1, \beta\delta, \beta\delta^2, \dots$ ⁶ Such naive agents believe that their future selves will make choices that are consistent with their current preferences. We adopt such naive beliefs because they increase the force of procrastination, but our qualitative results would be unchanged if we instead assumed that agents are sophisticated in their beliefs. For simplicity and analytical tractability, we set $\delta = 1$ (no long-run discounting).⁷ We also

⁵Another natural generalization is to consider the case in which agents have imperfect information about their personal value of s^* . If agents learn more about this value over time, they have another motive for delaying the costly action of opting out of the default.

⁶See Laibson (1997) for a discussion of hyperbolic discount functions and Akerlof (1992) and O'Donoghue and Rabin (1999) for a discussion of naifs and procrastination. Note that the term “hyperbolic” is overly restrictive, since the important property of these preferences is simply that they are characterized by more discounting in the short-run than in the long-run.

⁷We will calibrate our model at the frequency of a pay-cycle. So if the annual long-run discount rate is 0.05, then the discount rate per pay-cycle is approximately $0.05/26 = 0.002$ or $0.05/12 = 0.004$, implying respective δ values of 0.998 and 0.996. Relative to these values, setting $\delta = 1$ has little impact on our results.

adopt the standard hyperbolic assumption of $\beta < 1$.

We use the following timing convention. If the employee has not previously opted out of the default, the period begins with a flow loss of L . The employee then draws a current opt-out cost c and decides whether to delay opting out or to instead pay the cost, thereby ending the game. If the employee delays she will pay a flow cost of L next period and also face an anticipated continuation value function, which we denote $v(c')$, where c' represents next period's draw from the cost distribution. Hence, the employee chooses to pay c and end the game if the cost today is less than the discounted cost of delay, or $c < \beta [L + Ev(c')]$. When this inequality is not satisfied, the employee chooses to delay. Ignoring mixed strategies, which only arise on a zero measure region of the parameter space, the employee's strategy is thus

$$\begin{aligned}
 \text{"Opt out only when } c = 0\text{"} & \quad \text{if } \beta [L + Ev(c')] < 1 \\
 \text{"Opt out when } c = 0 \text{ or } c = 1\text{"} & \quad \text{if } \beta [L + Ev(c')] \geq 1
 \end{aligned} \tag{1}$$

1.1. Naive expectations and the continuation value function $v(c)$. Since the employee is assumed to be a naive hyperbolic agent, the continuation value function is constructed under the (mistaken) belief that all future selves will exhibit no time discounting, since this is what today's self wants those future selves to do. Recall that $\delta = 1$.

The strategy of opting out whatever the draw from the cost distribution means that the employee's expected loss is $\mu = E(c)$. Waiting until $c = 0$ to opt out implies that

the employee's expected loss would be

$$\begin{aligned} Ev(c|\text{wait until } c = 0) &= \mu [L + Ev(c|\text{wait until } c = 0)] \\ &= \frac{\mu L}{1 - \mu} \end{aligned}$$

This formula has a natural interpretation: the expected costs are equal to the expected per-period loss, μL , multiplied by the expected duration of the losses, $\frac{1}{1-\mu}$.

If $L < 1 - \mu$, then $\frac{\mu L}{1-\mu} < \mu$, implying that the expected losses generated by waiting to opt out until $c = 0$ are less than the losses from opting out immediately at cost $c = 1$. So if $L < 1 - \mu$, the employee will plan to wait until $c = 0$ to opt out. If $L \geq 1 - \mu$, the employee anticipates that next period she will act with certainty. In summary,

$$Ev(c) = \begin{cases} \frac{\mu L}{1-\mu} & \text{if } L < 1 - \mu \\ \mu & \text{if } L \geq 1 - \mu \end{cases} \quad (2)$$

We reiterate that $Ev(c)$ is based on naive beliefs, so this expectation reflects the actor's incorrect model of her future behavior.

1.2. Actual actions and welfare. Using (1) and (2), the probability of opting out in any period will be

$$p = \begin{cases} 1 - \mu & \text{if } L < \frac{1}{\beta} - \mu \\ 1 & \text{if } \frac{1}{\beta} - \mu \leq L \end{cases}$$

So the expected cost of opting out, conditional on opting out, will be

$$E(c|\text{opt out}) = \begin{cases} 0 & \text{if } L < \frac{1}{\beta} - \mu \\ \mu & \text{if } \frac{1}{\beta} - \mu \leq L \end{cases}$$

Let $w(c)$ represent the employee's expected total costs, discounted with the agent's long-run discount factor. A recursive representation for $w(c)$ is given by

$$\begin{aligned} Ew(c) &= pE(c|\text{opt out}) + (1-p)\delta[L + Ew(c')] \\ &= pE(c|\text{opt out}) + (1-p)[L + Ew(c')]. \end{aligned}$$

We evaluate social welfare using the long-run discount factor δ and omitting the short-run discount factor β . These preferences represent the actor's preferences at economic birth, which we assume occurs before she starts working at the firm. The last equation contains *no discounting*, since it reflects the fact that $\delta = 1$ in our calibration. Note however that our results would not change qualitatively if we had instead assumed $\delta < 1$ throughout our analysis.

Because $Ew(c) = Ew(c')$, we can show that

$$Ew(c) = \begin{cases} \frac{\mu L}{1-\mu} & \text{if } L < \frac{1}{\beta} - \mu \\ \mu & \text{if } \frac{1}{\beta} - \mu \leq L \end{cases}$$

We are now in a position to characterize the relationship between defaults and

welfare. To do this, we consider the relationship between expected (dis)utility and L , the per-period flow losses of not being at an optimum. To focus on the role of L , we stop suppressing L in our notation and consider

$$W(L) = Ew(c)|_L.$$

$W(L)$ is the expected losses for an agent with initial flow losses per period of L .

In a standard model with exponential discounting (i.e., $\beta = 1$), $W(L)$ would increase as flow costs L increase. But for hyperbolics (i.e., $\beta < 1$), it will always be the case that W is non-monotonic in L . To see this, note that $W(L) = \mu$ when $L = 1 - \mu$. This is the level of L at which an exponential (i.e., dynamically consistent) agent should opt out of the default whatever the cost realization. But when $c = 1$, a hyperbolic agent will only opt out of the default if $L \geq \frac{1}{\beta} - \mu$, which is greater than $1 - \mu$. Hence, when $1 - \mu < L < \frac{1}{\beta} - \mu$, the hyperbolic agent is insufficiently motivated to act, and this motivational gap produces self-defeating procrastination. In this region of L values, the expected loss function lies above μ , the value that $W(L)$ would take if the agent were not procrastinating and were willing to act at the high cost realization. But once L is high enough — specifically, above $\frac{1}{\beta} - \mu$ — the procrastination effect vanishes and expected costs fall back to μ , since the hyperbolic agent is now willing to act whatever the cost realization. Figure 1 plots the expected cost function against the flow costs L , revealing the non-monotonicity that arises whenever $\beta < 1$.

In a world with procrastination, moving the agent further from the optimum (i.e.,

increasing flow costs L) can make an agent *better* off, since it decreases the agent's tendency to procrastinate. This effect is not everywhere offset by the direct effect of reduced welfare arising from the increase in the delay cost, L .

1.3. The firm's optimization problem. We now analyze the employer's choice of a default savings rate under the assumption that the employer is interested in maximizing the welfare of the firm's employees. We recognize, however, that employer and employee incentives need not generally be aligned. This is particularly likely in the case presented here, since naive hyperbolic agents will not anticipate their own tendency to procrastinate and hence will not pick an employer based on the employer's ability to mitigate the harms of such procrastination. Therefore, this normative exercise is also relevant for regulators or unions that can influence the defaults that firms pick. Identifying and incorporating the other motivations and constraints that firms face in designing their benefit plans (e.g. non-discrimination testing, good corporate citizenship, reputational value in the labor market, or personal altruism, to name a few) is beyond the scope of the current paper.

We derive the optimal default, s_D^* , that minimizes the social welfare function,

$$\int_{\underline{s}}^{\bar{s}} W(L(s_D, s^*))f(s^*)ds^*. \quad (3)$$

We adopt the cost function

$$L(s_D, s^*) = \kappa(s_D - s^*)^2.$$

This quadratic cost function is convex in deviations from the optimal savings rate, s^* , and has the advantage of analytic tractability. However, it does not reflect the particular institutional features of many 401(k) plans (e.g., an employer match that ends at a threshold, implying a discontinuity in the cost function). We believe that the quadratic cost function represents a good compromise between tractability and realism.

We will minimize equation (3) numerically, using the actual estimated distribution of optimal savings rates. However, for the purposes of exposition, it is useful to consider the case in which $f(s^*)$ is uniform over support $[\underline{s}, \bar{s}]$. In this case, one can prove the following result when $\beta < 1$.

$$s_D^* = \begin{cases} \frac{\bar{s} + \underline{s}}{2} & \text{if } \bar{s} - \underline{s} \text{ small} \\ \underline{s} + \sqrt{\frac{1}{\kappa}(1 - \mu)} \text{ or } \bar{s} - \sqrt{\frac{1}{\kappa}(1 - \mu)} & \text{if } \bar{s} - \underline{s} \text{ large} \end{cases}$$

Intuitively, when there is little variation in optimal savings rates, it is best to design a default that is in the middle of the range of optimal savings rates, since all employees will then be very close to their optimal savings rate and delays in opting out of the default will not be very costly. By contrast, when there is a great deal of variation in optimal savings rates, it is better to design a default that is close to one of the

two boundaries of the support. This “boundary” strategy reduces the proportion of employees who engage in costly procrastination, since the boundary strategy reduces the fraction of employees who fall in the “procrastination” interval $1 - \mu < L < (1/\beta) - \mu$.

Finally, note that if $\beta = 1$ and f is uniform, then $s_D^* = \frac{\bar{s} + \underline{s}}{2}$ will *always* be an optimum⁸ because the procrastination effect does not apply and there is no gain in welfare from moving agents away from their optima.

It is also useful to emphasize a trivial property of these models, which is important in the empirical analysis that follows. This additional effect is easiest to understand if we assume that f is a discrete density on the domain of feasible savings rates: $\{0.00, 0.01, 0.02, \dots\}$. Then it is easy to show that

$$\lim_{\kappa \rightarrow \infty} s_D^* \in \arg \max_{s^*} f(s^*).$$

In other words, as the cost of deviations rises ($\kappa \rightarrow \infty$), the optimal default converges to the mode of the distribution of s^* . This effect is driven by the fact that for large costs of deviating from s^* , all employees will immediately adjust to their s^* except those who are already at their optima. Hence, the optimal social policy minimizes adjustment costs by setting the default equal to the most common value of s^* . We refer to this as the mode effect.

⁸However, it will not generally be the unique optimum.

1.4. Calibration. Our model has very few free parameters: the density of optimal savings rates, $f(s^*)$; the short-term discount factor β ; the scaling variable κ ; and the probability of a high-cost draw μ . We further restrict this list by using individual employee data to pin down the density f (see next section). We set $\beta = \frac{2}{3}$, reflecting a large body of experimental evidence and a growing body of field evidence. For example, Laibson, Repetto, and Tobacman (2003) use the method of simulated moments to estimate β using household financial data. Their benchmark estimate is 0.70 with a standard error of 0.11.

Only κ and μ remain to be calibrated. Before doing this we need to pick units for the variables in our model. We assume that time units are periods of a pay cycle (about two weeks). We assume that utility units can be interpreted in terms of a money metric in which one unit of utility is equal in value to 1/10 of a pay cycle of income. So when the cost realization is high ($c = 1$), opting out of the default generates a time cost that is equal in value to 1/10 of the agent's income during that pay cycle. We assume that such busyness is the norm and set $\mu = 0.9$. It then follows that the cost realization will be zero 1/10 $(1 - \mu)$ of the time.

To set κ , we use the following thought experiment. Suppose that a consumer is ten percentage points away from her optimal savings rate: $|s_D - s^*| = 0.1$. What is the money-metric cost of this deviation? Let x represent the loss in units of 1/10th of one pay cycle of income. Then, $\kappa(0.1)^2 = x$. We will consider a range of values for $x : 0.1, 1, 10$. This translates into the following range of values for $\kappa : 10, 100,$

1000. We consider this wide range for two reasons. First, we are agnostic about the appropriate calibration value. Second, we wish to explore the sensitivity of our results to the choice of κ . However, if forced to choose, we would set $\kappa = 100$, implying that a ten percentage point deviation in one's savings rate is as bad as losing one tenth of one's income during that pay cycle. For companies with an employer match, one could motivate losses of this magnitude by considering the missed match payments induced by undersaving.

2. EMPIRICAL ANALYSIS

Table 1 shows the variation in both 401(k) plan design and employee characteristics of the four companies for which we compute the optimal default 401(k) savings rate. We denote these four companies by their industry: Health, Office, Food, and Finance. All are large employers with well-established 401(k) plans.

There are two key differences in the 401(k) plan environment that vary across the companies. First, two of the companies (Health and Office) match employee contributions up to 6% of pay, while the other two have no match at all. These latter companies are of interest because the distribution of employee contribution rates will not be affected by the presence of a match threshold. Having an employer match may either raise or lower the desired 401(k) contribution rate. Because the match subsidizes saving in the 401(k) plan, employees with a match may desire to contribute more, at least up to the match threshold. However, the match also increases the total amount of savings that is being done, and the employees may use the match as a means to

offset their own contributions.

The second key difference in plan environment is that two of the companies (Office and Food) have an employer-sponsored defined benefit pension plan, while the other two do not. Other things equal, we would expect a lower desired savings rate for employees in companies with a defined benefit pension.

The workforce demographics of our four companies also vary quite considerably. The median pay ranges from \$25,000 per year in Food to \$41,000 per year in Finance. Because Social Security replaces a higher fraction of income for low income employees, we would expect a higher desired savings rate for high income employees. There is also significant variation in the fraction of employees that are female (from 30% in Office to 78% in Health) and the median age of the workforce (from 29 years in Finance to 39 years at Food).

To estimate the distribution of optimal savings rates (i.e., the density f in the model), we use two approaches. First, we report densities over 401(k) savings rates for “medium-tenure” employees. We informally reason that such medium-tenure employees have been at a firm long enough to select their optimal savings rate (i.e., the option value of waiting and procrastination hurdles have been surmounted), but not so long that tenure-driven selection effects dominate the data. These savings densities are reported in Table 2 for employees with 3-5 years of tenure (density f_1) and 5-7 years of tenure (density f_2).

Second, we use a regression framework to control for demographic variables. We

run an ordered logit regression in which the explanatory variable is the actual 401(k) contribution rate chosen by each individual employee. We include non-participation, which implies a 0% contribution rate, as one of the categories. The control variables in the regressions are $\ln(\text{pay})$, $\ln(\text{age})$, $\ln(\text{tenure})$, and a gender dummy variable ($D = 1$ if the employee is female). We then predict the distribution of contribution rates that would obtain if each employee had 30 years of tenure, holding other demographic characteristics constant. The underlying presumption behind this exercise is that 30 years is enough time to overcome any delays due to procrastination or the option value of waiting. The projected density from this procedure is reported as density f_3 in Table 2.

With these densities in hand, we are now in a position to estimate the optimal savings rate by minimizing equation (3), the social welfare function. We undertake this maximization for $3 \times 3 \times 4$ cases of interest: three different values for κ , three different ways of calculating the density f , and four different test companies. The results of these maximizations are reported in Table 3.

Table 3 documents six findings. First, the analysis reveals a high degree of heterogeneity in policy recommendations. The optimal default ranges from 0% to 15%. Moreover, even within a single firm there exists a large degree of variation in optimal defaults (e.g., Finance). Second, the range of variation in optimal defaults is twice as large as the range of average optimal savings rates. Third, the optimal default calculation is extremely sensitive to distributional assumptions on s^* . To see this, fix

$\kappa = 100$ and read across the columns. The defaults show substantial variation arising from very small (within-company) differences in f_1 , f_2 , and f_3 (see Table 2). Fourth, as κ gets large, much of the variation in optimal defaults is driven by the mode effect. For $\kappa = 1000$, five out of twelve of the optimal defaults are equal to the modal optimal savings rate. Fifth, the optimal defaults vary in a sensible way with the underlying firm-specific attributes. Firms whose employees have a high motive to save turn out to have higher optimal defaults than firms whose employees have a low motive to save. For example, the employees at Food have a defined benefit plan and a low average salary (i.e., a high average Social Security income replacement rate), and hence very low optimal defaults (0% to 3%). By contrast, the employees at Finance have no defined benefit plan, a high average salary, and a median optimal default of 14%. Sixth, and finally, the optimal defaults tend to cluster in one of three regions: close to 0%, close to the match threshold (6% for Health and Office), or close to the maximum contribution rate allowed under the plan.

3. DISCUSSION AND CONCLUSION

This paper has presented a model of 401(k) enrollment. The model includes four components: costs of opting out of a default, an option value of waiting to incur those costs, procrastination in opting out of a default, and heterogeneity in optimal savings rates.

One should also consider other important psychological and economic issues when picking socially optimal defaults. First, some employees may interpret defaults as im-

PLICIT advice, an issue that does not arise in the current model since each employee is assumed to know her true optimal savings rate.⁹ Second, defaults may be particularly sticky because of loss aversion.¹⁰ If the default is perceived to be a reference point, then deviations from that reference point may be psychologically aversive, since the resulting “gains” from the deviation (e.g., higher current consumption) are only weighted half as much as the resulting “losses” (e.g., lower saving). Third, if households do not know how to think about the future or are overoptimistic about future income, they may undervalue savings. In such a world, it may be optimal to pick a high default savings rate, even if households eventually move away from it. Fourth, households may know the optimal savings rate but not appreciate how important it is to implement it, increasing action delays. Fifth, choosing a long-run savings rate that is one percentage point too low is more costly than choosing a long-run savings rate that is one percentage point too high (since retirement is short relative to working life and the utility function generates a precautionary savings motive¹¹), suggesting a desirable upward shading of optimal defaults. Sixth, optimal savings rates are not constant over time (as we assume), but instead are likely to trend up slowly with working age. Seventh, the firm may wish to pick an optimal default that weights some employees more heavily than others. For example, it may be sensible to calculate optimal defaults that overweight the interests of employees that are likely to have a long duration of em-

⁹Employees may treat a zero default as weaker implicit advice than a non-zero default.

¹⁰See Kahneman and Tversky (1979), Thaler (1980), and Samuelson and Zeckhauser (1988) for a discussion of loss aversion and status quo bias.

¹¹Precautionary savings effects arise when $u''' > 0$, a common assumption in applied economic models.

ployment at the firm and underweight employees that are likely to separate relatively quickly. Future work should extend our theoretical framework by incorporating many of these additional considerations.

Future work should also explore the empirical implications of our model. The model makes quantitative predictions about the timing of savings rate changes. Employees who change their savings rate soon after they are hired should select larger changes than employees who change their savings rate long after they are hired. This is because employees who are willing to wait a long time for a low cost opportunity to opt out of the default are likely to have little to gain from doing so. The model also predicts that average savings rates will not necessarily increase monotonically with the default saving rate. As the default savings rate rises, procrastination effects can strengthen, leading more agents to delay selecting an even higher savings rate. Such perverse effects have already been observed in the data (Madrian and Shea 2001, Choi et al. 2003b).

Finally, the model suggests one important generalization that we are currently exploring (Choi et al. 2002b). If it is occasionally optimal to select “bad” defaults — i.e., defaults that are not close to one’s optimum saving rate — then it may be optimal to pick defaults that are so “bad” that all consumers feel compelled to immediately opt out of them. Such a setup is equivalent in practice to something that we call “active decision,” a regime that forces new employees to pick their own savings rate early in their tenure at the company without the benefit of a fall-back default. In a world with significant procrastination, such active decision regimes are sometimes the best

“defaults” of all.

4. REFERENCES

Akerlof, George A. “Procrastination and Obedience.” *American Economic Review*, May 1991 (*Papers and Proceedings*), 81(2), pp. 1-19.

Choi, James J.; Laibson, David; Madrian, Brigitte and Metrick, Andrew. “Defined Contribution Pensions: Plan Rules, Participant Decisions, and the Path of Least Resistance,” in James M. Poterba, ed., *Tax Policy and the Economy*, Vol. 16. Cambridge, MA: MIT Press, 2002a, pp. 67-113.

_____. “Active Decisions: A Natural Experiment in Saving.” Mimeo, Harvard University, 2002b.

_____. “Optimal Defaults.” *American Economic Review*, May 2003a (*Papers and Proceedings*), 93(2), pp. 180-185.

_____. “For Better or For Worse: Default Effects and 401(k) Savings Behavior,” in David Wise, ed., *Perspectives in the Economics of Aging*. Chicago: University of Chicago Press, 2003b.

Kahneman, Daniel and Tversky, Amos. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica*, March 1979, 47(2), pp. 263-292.

Laibson, David. “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics*, May 1997, 112(2), pp. 443-77.

- Laibson, David; Repetto, Andrea and Tobacman, Jeremy.** “Instant Gratification Over the Lifecycle.” Mimeo, Harvard University, 2003.
- Madrian, Brigitte C. and Shea, Dennis.** “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior.” *Quarterly Journal of Economics*, November 2001, *116*(4), pp. 1149-1187.
- O’Donoghue, Ted and Rabin, Matthew.** “Doing It Now or Later.” *American Economic Review*, March 1999, *89*(1), pp. 103-24.
- Samuelson, William and Zeckhauser, Richard.** “Status Quo Bias in Decision Making.” *Journal of Risk and Uncertainty*, March 1988, *1*(1), pp. 7-59.
- Sunstein, Cass, and Thaler, Richard H.** “Paternalism.” *American Economic Review*, May 2003a (*Papers and Proceedings*), *93*(2), pp. 175-179.
- Thaler, Richard H.** “Toward a Positive Theory of Consumer Choice.” *Journal of Economic Behavior and Organization*, March 1980, *1*(1), pp. 39-60.

Figure 1: Expected total losses as a function of flow cost per period

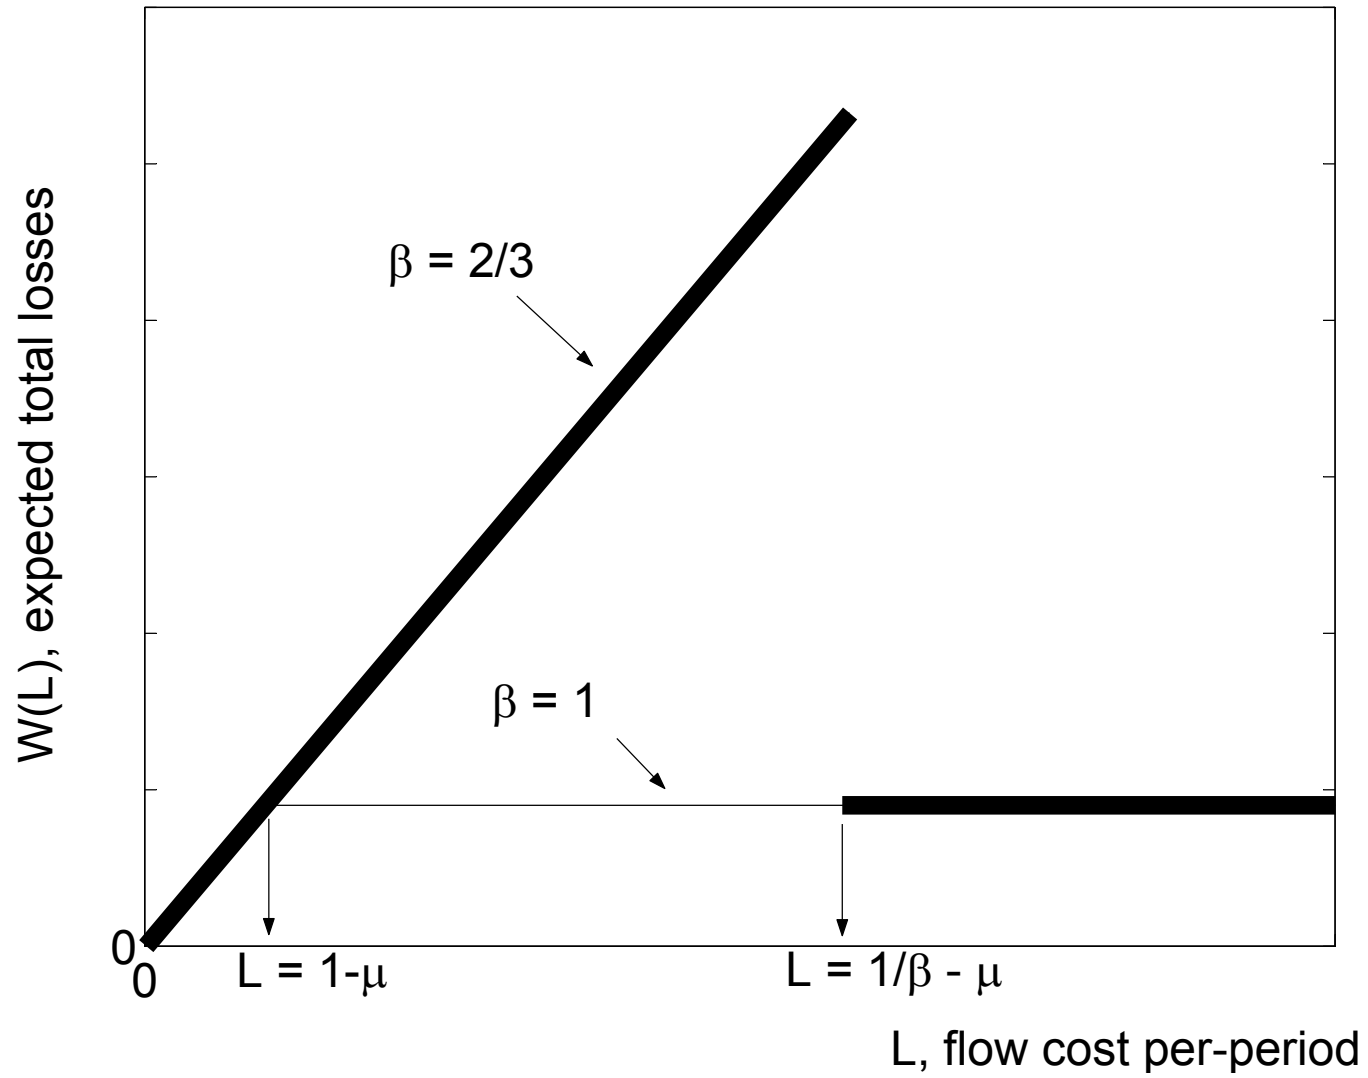


Table 1. Characteristics of Employees and Their 401(k) Plans

	Health Company	Office Company	Food Company	Financial Company
Employer Match	\$0.50/\$1 to 6%	\$0.67/\$1 to 6%	None	None
Contribution rate range	0% to 15%	0% to 16%	0% to 15%	0% to 15%
Company DB plan	No	Yes	Yes	No
401(k) participation rate	61.9%	74.2%	32.8%	63.4%
Avg. 401(k) contribution rate	4.3%	4.5%	2.0%	6.0%
Median salary	\$31,034	\$27,629	\$25,355	\$41,109
Median age	37.9 years	36.7 years	38.5 years	28.9 years
Median tenure	4.8 years	5.4 years	5.6 years	2.0 years
Fraction female	77.8%	30.1%	54.0%	50.0%
Year	1997	1998	1998	1998

Source: Company summary plan descriptions and calculations of the authors. The sample in column 1 is all employees with 1+ year of tenure. The sample in column 2 is all employees with 2+ years of tenure. The sample in columns 3 and 4 is all employees.

Table 2. Savings Rate Distributions

This table reports distributions of savings rates. $f_1(s^*)$ is the savings rate distribution of eligible employees at December 31, 1997 whose tenure is between 3 and 5 years. $f_2(s^*)$ is the savings rate distribution of eligible employees at December 31, 1997 whose tenure is between 5 and 7 years. $f_3(s^*)$ is the distribution of optimal savings rates based on predicted values from an ordered logit regression of savings rate on age, gender, pay, and tenure. Predicted values are calculated using 30 years of tenure instead of actual tenure.

s^*	Health Company			Office Company			Food Company			Financial Company		
	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$
0%	0.35	0.26	0.20	0.32	0.23	0.12	0.72	0.59	0.51	0.22	0.19	0.19
1%	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.04	0.04	0.02	0.02	0.01
2%	0.06	0.06	0.03	0.15	0.14	0.09	0.04	0.06	0.06	0.02	0.02	0.02
3%	0.06	0.07	0.05	0.05	0.05	0.04	0.04	0.05	0.05	0.04	0.02	0.02
4%	0.04	0.05	0.03	0.04	0.05	0.04	0.02	0.03	0.03	0.03	0.03	0.02
5%	0.04	0.05	0.04	0.04	0.04	0.04	0.06	0.09	0.12	0.09	0.09	0.07
6%	0.26	0.28	0.30	0.24	0.31	0.36	0.01	0.02	0.02	0.04	0.05	0.04
7%	0.02	0.02	0.02	0.01	0.01	0.02	0.01	0.01	0.01	0.04	0.04	0.03
8%	0.02	0.03	0.03	0.02	0.02	0.03	0.01	0.01	0.01	0.05	0.04	0.04
9%	0.01	0.01	0.01	0.00	0.00	0.01	0.00	0.00	0.00	0.02	0.02	0.02
10%	0.05	0.06	0.09	0.06	0.06	0.10	0.03	0.04	0.06	0.13	0.13	0.13
11%	0.01	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.02	0.02
12%	0.01	0.01	0.02	0.01	0.01	0.02	0.00	0.00	0.00	0.03	0.03	0.03
13%	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01
14%	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06	0.08	0.06
15%	0.07	0.09	0.15	0.01	0.01	0.01	0.04	0.05	0.08	0.20	0.20	0.30
16%	NA	NA	NA	0.04	0.05	0.11	NA	NA	NA	NA	NA	NA
Mean	4.29%	5.00%	6.40%	4.02%	4.65%	6.43%	1.77%	2.40%	3.19%	7.50%	7.97%	8.82%

Table 3. Optimal Default Savings Rates

This table shows the optimal savings rate for four different firms. Food Company and Financial Company have no employer match in their plans. $f_1(s^*)$ is the savings rate distribution of eligible employees at December 31, 1997 whose tenure is between 3 and 5 years. $f_2(s^*)$ is the savings rate distribution of eligible employees at December 31, 1997 whose tenure is between 5 and 7 years. $f_3(s^*)$ is the distribution of optimal savings rates based on predicted values from an ordered logit regression of savings rate on age, gender, pay, and tenure. Predicted values are calculated using 30 years of tenure instead of actual tenure.

κ	Health Company			Office Company			Food Company			Financial Company		
	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$	$f_1(s^*)$	$f_2(s^*)$	$f_3(s^*)$
10	4%	5%	6%	4%	5%	6%	2%	2%	3%	7%	8%	9%
100	2%	2%	14%	2%	2%	5%	1%	1%	2%	2%	14%	14%
1000	0%	0%	15%	6%	6%	6%	0%	0%	0%	15%	15%	15%
Mean	4.29%	5.00%	6.40%	4.02%	4.65%	6.43%	1.77%	2.40%	3.19%	7.50%	7.97%	8.82%
Mode	0.00%	6.00%	6.00%	0.00%	6.00%	6.00%	0.00%	0.00%	0.00%	0.00%	15.00%	15.00%