The Limitations of Defaults

John Beshears
Stanford University and NBER

James J. Choi
Yale University and NBER

David Laibson
Harvard University and NBER

Brigitte C. Madrian
Harvard University and NBER

Prepared for the 12th Annual Joint Conference of the Retirement Research Consortium

August 5-6, 2010
Washington, D.C.

Revised September 15, 2010

The research reported herein was supported by the U.S. Social Security Administration (SSA) through grant #10-M-98363-1-02 to the National Bureau of Economic Research (NBER) as part of the Retirement Research Consortium (RRC). The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, the RRC, or NBER. The authors thank Yeguang Chi for excellent research assistance and also acknowledge individual and collective financial support from the National Institute on Aging (grants R01-AG-021650 and T32-AG-000186).

Abstract

Prior research has demonstrated that defaults have a powerful influence on economic outcomes in a wide range of settings because individuals often passively accept default options. This paper examines the degree to which defaults become less powerful as they become more extreme. We study a firm with a defined contribution retirement savings plan in which employees are automatically enrolled at a 12% contribution rate, a rate that is considerably higher than those studied in previous work. In addition, the default contribution rate is suboptimal for all employees because the firm only matches employee contributions between 12% and 18% of pay. Approximately one-quarter of employees at this firm remain at the default contribution rate after twelve months of tenure, while the comparable fraction for firms with more modest defaults is more than 60%. We also find that employees who remain at the default contribution rate after twelve months of tenure have lower incomes than would be predicted by the incomes of employees who actively choose neighboring contribution rates. This evidence suggests that defaults are more influential for low-income employees than for high-income employees because low-income individuals generally face higher barriers to active decision-making.

A wide range of evidence has documented that economic outcomes can be powerfully influenced by default options – the options that are selected on behalf of an individual when the individual does not affirmatively make selections herself. In decision domains as diverse as e-mail marketing (Johnson, Bellman, and Lohse, 2003), organ donation (Johnson and Goldstein, 2003; Abadie and Gay, 2004), automobile purchases (Park, Jun, and MacInnis, 2000), and retirement savings and investment (Madrian and Shea, 2001; Choi et al., 2002 and 2004; Beshears et al., 2008), individuals are likely to accept the default passively, so the policymaker or manager who chooses the default is often able to choose the modal outcome that is implemented.

In this paper, however, we study the limitations of defaults. Given the effectiveness of defaults in a variety of settings, it is natural to ask how far the influence of defaults can extend: to what extent will individuals more actively resist a default when it begins to conflict with their best interests? To better understand the limitations of defaults, we also examine the closely related issue of which employees are more likely to be influenced by defaults. Previous research has found that low-income individuals are less likely to opt out of defaults (Choi et al., 2004). Here, we ask whether this pattern is primarily the result of defaults implementing outcomes close to those that low-income individuals would have selected for themselves had they made active choices, or primarily the result of low-income individuals facing higher barriers to active decision-making, such as cognitive costs, costs of action, a tendency to procrastinate, or a lack of information or expertise.

To shed light on these questions, we analyze data on the defined contribution retirement savings plan of a firm in the United Kingdom. Eligible employees at the firm are automatically enrolled in the plan upon hire at a 12% default contribution rate, with contributions invested in a default asset allocation that is a mix of bonds and equities. In other words, an employee who does not actively elect otherwise will contribute 12% of every paycheck to the plan on a before-tax basis, invested according to the default asset allocation. This default contribution rate is considerably higher than the default contribution rates that have been examined in the past, which are typically in the range of 2% or 3% and sometimes as high as 6% (Choi et al., 2002 and 2004).

The default contribution rate at the firm we study is not only a particularly aggressive default option but also a suboptimal choice for all employees. The firm contributes to employee

accounts by matching employee contributions between 12% and 18% of pay on a one-for-one basis. In a stylized two-period model where the employee divides resources between present and future consumption (see Figure 1), this employer match structure creates a non-convexity in the employee's budget set. A standard indifference curve cannot be tangential to the budget set at the point corresponding to a 12% contribution rate, where there is a kink. Furthermore, if an employee would choose a contribution rate close to 12% in the two-period context, the same employee in a multi-period context would be better served by switching back and forth between a lower contribution rate (below 12%) and a higher contribution rate (above 12%), since this strategy convexifies the budget set. Thus, the default contribution rate in the plan should not be an attractive option for any employee.[1]

Using data on employees hired at the firm between July 2006 and June 2007, we analyze the extent to which employees opt out of this unattractive default. By twelve months of tenure, only 25% of employees have passively stayed at the 12% default contribution rate. At two firms studied by Choi et al. (2004) with more modest default contribution rates (2% and 3%), the comparable fractions were 61% and 66%.

Opt-out behavior along the asset allocation dimension is strikingly different. In the retirement savings plan we study, 66% of employees remain at the default investment allocation for their first twelve months of tenure, even though 73% of those who remain at the default investment allocation have opted out of the default contribution rate. This pattern suggests that the high opt-out rate from the contribution default is not driven by the employee population being particularly intensely engaged in the details of their financial affairs. The pattern is consistent with the hypothesis that employees have some sense of their optimal contribution rates but little expertise when approaching the multi-dimensional asset allocation problem, making them more likely to rely on the default asset allocation for guidance. Of course, the evidence is also consistent with the possibility that the default asset allocation in the plan is close to the optimum for many employees.

[1] As we explain below in Section I, an employee at the firm who elects a contribution rate higher than 12% in order to earn matching contributions must in general maintain that contribution rate until the next annual open enrollment period. If an employee plans to leave after only a brief tenure with the firm, the dynamic strategy of switching between low and high contribution rates may not be feasible, and 12% may in fact be the employee's optimal contribution rate because of integer constraints. However, while this scenario is logically possible, we suspect that it is rare, and our discussion presumes that it does not apply to any employee.

The evidence on opt-out behavior suggests that defaults can quickly lose much of their power as they become more aggressive, but in the case we study, a meaningful fraction of individuals is still slow to opt out of the default. We find that the employees at the company who remain at the 12% default contribution rate after twelve months of tenure have incomes that are approximately one-third lower than the level that is predicted from a regression of income on contribution rate among employees who actively choose to opt out of the default. This result echoes previous work documenting that low-income individuals are slower to opt out of automatic enrollment defaults than high-income individuals (Choi et al., 2004). However, those previous findings are based on retirement savings plans with low default contribution rates (e.g., 3%), making it unclear whether low-income employees are slower to opt out largely because the default is close to the options that they would have chosen for themselves anyway, or because they face higher barriers to active decision-making, such as procrastination or a lack of expertise. The savings plan we study in this paper gives us some ability to distinguish between these two explanations. Because the default contribution rate in this plan is 12%, it is likely that low-income employees have stronger financial reasons to opt out to a different (lower) contribution rate than high-income employees, so the fact that low-income employees are still slower to opt out suggests that barriers to active decision-making are more important for these employees. Policymakers and managers may wish to keep in mind that defaults are likely to affect this group of people disproportionately.

The paper proceeds as follows. In Section I, we explain the structure of the retirement savings plan. In Section II, we describe our data. We document default opt-out patterns in Section III and analyze differences in opt-out behavior between low-income and high-income employees in Section IV. We offer concluding remarks in Section V.

I. Defined Contribution Plan Design

The company we study has more than 50,000 employees engaged in a range of job functions, including manufacturing, marketing, research and development, and administration. We study the firm during the period July 2006 through June 2008. The financial crisis did not have a differentially adverse impact on the firm relative to an average firm.

The firm has its headquarters in the United Kingdom, and we limit our attention to the pension plan for U.K. employees. The U.K. pension system, which is broadly similar to that in

the United States, consists of three tiers. The first tier, the Basic State Pension, is a mandatory government scheme to which individuals contribute throughout their working lives in return for an annuity stream in retirement.[2] The second tier, the State Second Pension, is also a government scheme, but it is less progressive in the provision of benefits, as payouts in retirement are more closely linked to lifetime earnings.[3] The third tier is the system of private retirement savings plans. Contributions to these plans are tax-deductible for individuals (up to a limit[4]) and are generally tax-deductible for employers.[5]

In 2006, slightly more than half of U.K. workers were enrolled in a private retirement savings plan, and of these workers approximately one-third had a defined contribution plan.[6] The company that we study maintains legacy defined benefit plans for some of its employees, but all employees hired during the time period we study are eligible only for a defined contribution plan. We restrict our analysis to the firm's main defined contribution plan. Less than one percent of employees hired during this period are not eligible for the main plan but are instead eligible for a plan with a different structure.[7] These employees generally have low salaries, and we exclude them from our analysis because they face distinct plan rules and because they are too few in number to be examined separately.

All new employees besides the small group described above are eligible for the firm's main defined contribution plan and are automatically enrolled upon hire at a 12% default contribution rate. Employees can opt out of the plan entirely, but in order to remain active plan participants they must contribute at least 4% of every paycheck to the plan.[8] Subject to the 4% floor, employees can elect any contribution rate at any time.[9] The firm offers one-for-one

[2] In 2009, a complete contribution record would entitle an individual to £95.25 per week from the Basic State Pension.
[3] Both the first tier and the second tier are "pay-as-you-go" schemes. It is possible for workers to "contract out" of the second tier by contributing to a private pension instead of the State Second Pension.
[4] The 2009-2010 annual limit on tax-deductible contributions for individuals was the lesser of £245,000 and 100% of annual income. A lifetime limit also applies.
[5] The information in this paragraph is from the Pensions Policy Institute (2010).
[6] These figures are derived from data from the Office for National Statistics (2008). Public sector workers, who almost always have defined benefit plans, are included in the sample. Their employer-sponsored plans are considered "private" in this context simply to denote that the plans are distinct from the Basic State Pension and the State Second Pension.
[7] These employees are automatically enrolled at a default contribution rate of 4%, and they begin to receive an employer match when their contribution rate exceeds 4%.
[8] The firm occasionally allows an employee to remain a plan participant while contributing less than 4% of pay, but this privilege is granted on a case-by-case basis.
[9] Some fraction of the first 12% of employee contributions is actually designated as employer contributions for the purposes of determining National Insurance contribution levels. We do not observe the magnitude of the fraction.

employer matching contributions, which vest immediately, on employee contributions between 12% and 18% of pay. However, in order to obtain matching contributions, an employee must elect a contribution rate greater than 12% within her first three months of hire or within the three-week open enrollment period in late May and early June, and the employee must agree to maintain her selected contribution rate until the next open enrollment period. Employees who choose contribution rates greater than 12% outside of the designated windows do not receive matching contributions.[10] All contributions to the plan are before-tax, and loans from the plan are not permitted.

Plan account balances can be allocated according to the employee's wishes across eleven different investment funds, one of which is a cash fund, two of which are bond funds, and eight of which are equity funds. During the time period studied, the plan did not offer target retirement date funds (which slowly shift their asset mixes from equities to fixed-income investments over time) or employer stock in the investment menu. Employees who do not elect otherwise have their contributions invested according to a default asset allocation, which features a mix of bonds and equities.

II. Data on Defined Contribution Plan Outcomes

Our analysis of the company retirement plan relies on monthly administrative plan records from three data extracts. The first extract covers March 2006 through October 2007; the second extract covers November 2007 through March 2008; and the third extract covers April 2008 through June 2008. Each extract includes all employees who were active participants in the plan as of the end of the extract period (October 31, 2007; March 31, 2008; or June 30, 2008). We restrict our attention to the 671 employees who began their tenure at the firm between July 1, 2006, and July 1, 2007,[11] and who have data records for their first twelve full months of

---

The designation affects neither the amount of money that is credited to employee defined contribution accounts nor the corresponding deduction from employee pay, but the designation does reduce payments to the National Insurance system. Despite the relabeling of this portion of contributions, we still refer to the contributions as "employee contributions" because this term most accurately reflects the relationship between pay deductions and cash flows into employee accounts.

[10] In rare individual cases, the firm can allow an employee to earn matching contributions by choosing a contribution rate greater than 12% outside the designated windows, to change a match-earning contribution rate before the next open enrollment period, or to earn matching contributions even with a contribution rate less than or equal to 12%.

[11] We do not include employees hired between March and June of 2006 because the retirement plan rules were in flux during that period.

employment.[12] Note that our analysis excludes employees who left the firm or plan before the end of their twelfth tenure month and excludes employees who left the firm or plan after the end of their twelfth tenure month but before the end of the extract period that would have included their twelfth tenure month.

The data set includes the gender, marital status, age, and hire date of each employee. In addition, for each month, we observe employee compensation, the value of employee contributions to the plan, and the value of employer contributions to the plan. To calculate employee and employer contribution rates, we divide contributions by compensation. However, we make some adjustments to these calculations because administrative processes in the retirement savings plan often lag those in the employee payroll system. For instance, if an employee receives a pay raise, the compensation record will increase accordingly, but the plan contributions may stay at the amount appropriate for the previous compensation level, leading to a misleadingly low ratio of contributions to compensation. In this example, the subsequent month's contribution amounts may (or may not) adjust upwards to reflect the new compensation level *and* to make up for the missed contributions in the previous month, leading to a misleadingly high ratio of contributions to compensation. More complicated scenarios arise when an employee experiences multiple salary changes within a short timeframe. In addition, a similar issue affects plan contributions at the beginning of an employee's tenure: contributions in the first or second full tenure month may represent contributions for that month *and* for previous month(s). In all of these cases, we reattribute contributions to the appropriate months before generating contribution rates.

Another factor that affects the calculation of contribution rates is the ability of employees to contribute to the savings plan out of their bonus pay. Bonuses do not appear in our compensation data, so plan contributions out of bonuses can create misleadingly large contribution rates. Our analysis attempts to ignore contributions out of bonus compensation by adopting the following procedure. Because bonuses are often awarded in April, any employee contribution rate in April that exceeds the March contribution rate and the May contribution rate by more than six percentage points is set equal to the March contribution rate. After the contribution rate calculations have been corrected for misalignment between the compensation

[12] If an individual begins employment on the first working day of a month, that month is tenure month one. If an individual begins employment on a later day in the month, the subsequent month is tenure month one.

and contribution records and adjusted for contributions out of bonus pay, there still remain some non-integer contribution rates, and these are rounded to the nearest integer rate.

Our data set does not include a variable indicating which employees are participants in the firm's main defined contribution plan. However, the difference in structure between the main plan and the other plan allows us to identify employees who are likely to be members of the other plan. The main plan provides employer matching contributions when the employee contribution rate exceeds 12%, while the other plan provides a match when the employee contribution rate exceeds 4%. An employee who receives a match on contributions above 4% of pay would therefore be identified as a participant in the other plan, although no such employees in fact exist in our sample. To be conservative, our analysis excludes employees whom we never observe contributing more than 4% of pay, even though some of these individuals may be participants in the main plan. This restriction eliminates five employees, or 0.7% of the sample, a fraction that is in line with the fact that less than one percent of employees are eligible for the other plan instead of the main plan.

Finally, our data set includes information on employee asset allocations. On a monthly basis, we observe the value of shares bought or sold in each mutual fund in the investment menu, as well as variables indicating whether an employee has ever opted out of the default asset allocation for contribution flows and whether an employee has ever reallocated existing balances across funds.[13]

Table 1 presents summary statistics for our sample. More than half of the employees are female, and slightly less than half are married. The mean age is 35 years. At £28,700, the median salary is slightly higher than the typical salary of a U.K. worker, but there is considerable variation in pay across the firm's employees. The mean employee contribution rate at twelve months of tenure is 9.4% of pay, and the mean employer contribution rate is only 0.9% of pay, reflecting the fact that the firm does not match employee contributions until they exceed 12% of pay.

[13] In some cases, the variable for whether an employee has ever reallocated existing balances indicates that an employee has made such a change a few months before the change appears in the data on mutual fund flows. We rely on the mutual fund flow data when these discrepancies arise.

III. Default Opt-Out Patterns

In this section, we analyze the rate at which employees opt out of the savings plan defaults. We are particularly interested in opt-out behavior vis-à-vis the 12% default contribution rate. Previous studies of automatic enrollment in retirement savings plans have focused on firms that use default contribution rates between 2% and 6% of pay, so the 12% default can give us some insight into employee reactions to more aggressive defaults. Furthermore, as explained in Section I, the savings plan features employer matching contributions on employee contributions between 12% and 18% of pay, provided that the employee elects a contribution rate greater than 12% during a designated window of time and agrees to maintain that contribution rate. The budget set non-convexity created by this match structure implies that the 12% default contribution rate is not an optimal choice for any employee,[14] so opt-out behavior in this setting sheds light on how employees respond when a default is contrary to their best interests.

Figure 2 summarizes employee contribution rates at the firm as tenure increases. For each tenure month, the gray bar represents the fraction of employees in our sample who have never opted out of the 12% default contribution rate; the white bar represents the fraction who have chosen a contribution rate lower than 12%; the small black bar represents the fraction who have opted out of the default but returned to a 12% contribution rate; and the striped bar represents the fraction who have chosen a contribution rate higher than 12%. In these calculations, we disregard contributions out of bonus pay because they are infrequent occurrences that involve a separate decision-making process. The figure indicates that employees opt out of the default rapidly. By tenure month three, approximately one-third of the employees in our sample remain at the default, and this fraction steadily declines to one-quarter by tenure month twelve. As a point of contrast, consider two of the savings plans with automatic enrollment studied by Choi et al. (2004), which have default contribution rates of 2% and 3%. At these firms, 61% and 66% of employees with 12-17 months of tenure are contributing at the default rate.[15]

In Figure 2, it is apparent that nearly three-quarters of the employees who have opted out of the default contribution rate in the plan we study choose a rate lower than 12% in tenure month twelve. Figure 3 examines employee contribution rates at twelve months of tenure more

[14] Strictly speaking, there may be some employees for whom 12% is the optimal contribution rate, but this group of employees is probably small. See footnote 1.

[15] Note that the Choi et al. samples include employees who are not participating in the plan – they constitute 9% of the first sample and 13% of the second sample. The sample studied in the current paper excludes non-participants.

closely. As previous studies have documented (see, for example, Choi et al., 2004), many employees contribute at the minimum rate required to receive the full employer match – in this case, 10% of the sample has a contribution rate of 18%. However, 31% of the sample chooses a contribution rate of 4%, which is the lowest officially permissible rate for employees who wish to remain active plan participants (a small number of employees receive special permission to participate at a lower contribution rate). Thus, many employees who opt out of the default reject the 12% contribution rate decisively instead of adjusting their contribution rates incrementally.

Opt-out patterns on the asset allocation dimension are quite different from those on the contribution rate dimension. Figure 4 shows that approximately two-thirds of the sample has never opted out of the asset allocation default by tenure month twelve. This outcome is similar to the outcomes observed at the two automatic enrollment plans studied by Choi et al. (2004), who find that 57% and 64% of employees with 12-17 months of tenure have maintained the default asset allocation, which in the first case is 100% of assets invested in a stable value fund and in the second case is 100% of assets invested in a money market fund. For these calculations, an employee is deemed to have opted out of the asset allocation default if she either changes the allocation of her contribution flows or changes the allocation of existing plan balances.

Figure 5 combines information on contribution rate opt-out behavior with information on asset allocation opt-out behavior. At twelve months of tenure, 18% of the sample has never opted out of the contribution rate default or the asset allocation default, while 27% has opted out of both. Interestingly, 48% has opted out of the contribution rate default but not the asset allocation default, and the reverse is true for only 7% of the sample. It is possible that the asset allocation default has a greater impact on outcomes than the contribution rate default because individuals have some confidence in their ability to choose an appropriate savings rate but have little confidence in their ability to choose an appropriate asset allocation. Such individuals may opt out of the default contribution rate but maintain the default asset allocation, which they perceive as implicitly endorsed by their employer. Of course, it is also possible that many employees keep the default asset allocation because it is close to their optimal asset allocation.

IV. Opt-Out Patterns Among Low-Income and High-Income Employees

In the previous section, we presented evidence examining the extent to which individuals opt out of defaults as those defaults become more aggressive. In this section, we explore the related issue of which employees are more likely to opt out of defaults.

For Figure 6, we group employees in our sample into eight different categories based on their contribution rates at tenure month twelve. Employees with a contribution rate of 12% form one group, but other groups are based on pairs of contribution rates. For example, employees with contribution rates of 13% or 14% are grouped together. In this figure and in the regressions that accompany it (see Table 2), contribution rates less than 4% are recoded as being equal to 4%, and contribution rates greater than 18% are recoded as being equal to 18%, although the results are nearly identical if employees with contribution rates less than 4% or greater than 18% are dropped from the sample. For each group of employees, the box plotted in Figure 6 gives the mean of the logarithm of annual salary. Here, annual salary is simply the sum of monthly compensation over the first twelve full months of tenure. It is apparent from the figure that employees contributing to the savings plan at a 12% rate have lower salaries on average than employees who choose a slightly higher or lower contribution rate.

To formalize this difference in salaries, we perform an ordinary least squares regression of the logarithm of annual salary on the employee contribution rate, the employee contribution rate squared, and an indicator variable for the employee contribution rate being equal to 12%. The fitted values from this regression, restricting the indicator variable to be zero at all contribution rates, are given by the solid line in Figure 6, and the dotted lines represent two standard errors on either side of the fitted values. Employees with 12% contribution rates have salaries that are approximately one-third lower on average than the level we would predict based on the salaries of other employees, and this difference is highly statistically significant. Previous studies have documented that low-income employees are slower to opt out of defaults than high-income employees (Choi et al., 2004), and the findings from this savings plan are consistent with those results. In contrast with prior research, however, this savings plan helps us make inferences regarding the source of the differential propensity to opt out of defaults. Figure 6 provides evidence that the default contribution rate of 12% is far from the contribution rates that low-income employees would choose for themselves if they were to make active decisions, so the finding that low-income employees are nonetheless more likely to remain at the default than

high-income employees suggests that low-income employees experience higher barriers to active decision-making, such as procrastination or a lack of expertise. An important caveat to this analysis is that low-income employees may not have stronger financial incentives to opt out of the 12% default contribution rate, as high-income employees may have a greater capacity to opt out of the default and obtain employer matching contributions. This consideration may make high-income employees more motivated to deviate from the default than low-income employees.

The regression results used to construct Figure 6 are reported in column 3 of Table 2. This regression specification includes both the contribution rate and the squared contribution rate, so we test the robustness of our results by dropping the squared term (column 1) or by adding a cubed term (column 5). We also try fitting the data using a linear spline in the contribution rate with a knot point at 12% (column 7). All of these specifications give similar results. The coefficient on the indicator variable for having a contribution rate of 12% ranges from -0.348 to -0.363 and is always statistically significant at the 1% level. In columns 2, 4, 6, and 8 of Table 2, we add controls for gender, marital status, age, and month of hire to the regressions, and the results remain similar.

V. Conclusion

Participant behavior in the unique retirement savings plan we study can provide some insight into the limitations of defaults. The firm automatically enrolls employees at a 12% contribution rate, a default that is higher than the contribution rate defaults studied in the past and that is a suboptimal option for employees given the firm's offer to match employee contributions between 12% and 18% of pay. By tenure month twelve, three-quarters of employees have opted out of the default contribution rate, and many of these employees have chosen lower contribution rates. Thus, if it is socially desirable to have high contribution rates in defined contribution retirement savings plans, aggressive defaults may not be an effective policy – it may be more effective to automatically enroll employees at a more modest contribution rate (e.g., 6%) and to automatically increase contribution rates over time (see Thaler and Benartzi, 2004). Further research on this issue would be valuable.

On the question of which employees are more likely to remain at the default, we find that employees who remain at the 12% default contribution rate after twelve months of tenure have lower salaries on average than employees who elect nearby contribution rates. The evidence

suggests that barriers to active decision-making, such as procrastination or a lack of relevant knowledge, play an important role in the tendency of low-income employees to opt out of defaults more slowly than high-income employees. The particularly powerful impact of defaults on individuals with low human capital implies that variation in defaults across decision domains may cause these individuals to exhibit seemingly contradictory preferences. Policymakers may wish to keep these individuals in mind when setting defaults.

References

Abadie, Alberto, and Sebastien Gay, 2006. "The impact of presumed consent legislation on cadaveric organ donation: A cross-country study." *Journal of Health Economics* 25, pp. 599-620.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, 2008. "The importance of default options for retirement saving outcomes: Evidence from the United States." In Stephen J. Kay and Tapen Sinha, eds., *Lessons from Pension Reform in the Americas*, Oxford: Oxford University Press, pp. 59-87.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2002. "Defined contribution pensions: Plan rules, participant decisions, and the path of least resistance." In James Poterba, ed., *Tax Policy and the Economy*, Cambridge: MIT Press, pp. 67-113.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2004. "For better or for worse: Default effects and 401(k) savings behavior." In David A. Wise, ed., *Perspectives on the Economics of Aging*, Chicago: University of Chicago Press, pp. 81-121.

Johnson, Eric J., Steven Bellman, and Gerald L. Lohse, 2002. "Defaults, framing and privacy: Why opting in-opting out." *Marketing Letters* 13, pp. 5-15.

Johnson, Eric J., and Daniel Goldstein, 2003. "Do defaults save lives?" *Science* 302, pp. 1338-1339.

Madrian, Brigitte C., and Dennis F. Shea, 2001. "The power of suggestion: Inertia in 401(k) participation and savings behavior." *Quarterly Journal of Economics* 116, pp. 1149-1187.

Office for National Statistics, 2008. *Annual Survey of Hours and Earnings 2006*. Available at http://www.statistics.gov.uk/statbase/Product.asp?vlnk=14058, accessed May 1, 2010.

Park, C. Whan, Sung Youl Jun, and Deborah J. MacInnis, 2000. "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, pp. 187-202.

Pensions Policy Institute, 2010. *The Pensions Primer*. Available at http://www.pensionspolicyinstitute.org.uk/default.asp?p=97, accessed June 6, 2010.

Thaler, Richard H., and Shlomo Benartzi, 2004. “Save more tomorrow™: Using behavioral economics to increase employee saving.” *Journal of Political Economy* 112, pp. S164-S187.

Table 1. Sample Characteristics

This table presents summary statistics for the 671 employees whom we observe in the data for at least twelve months. The variables are measured as of tenure month twelve for each employee. For the employee contribution rate, contributions out of bonuses are disregarded.

	Mean	Std. Dev.	10^{th} Percentile	Median	90^{th} Percentile
Female	55.0%				
Married	47.7%				
Age (years)	35.0	9.4	24.4	33.1	48.5
Annual salary (£1000s)	35.3	22.4	15.7	28.7	64.5
Employee contribution rate (percent of pay)	9.4	5.7	4.0	9.0	18.0
Employer contribution rate (percent of pay)	0.9	2.0	0.0	0.0	6.0

Table 2. Regression Analysis of Employee Salaries by Contribution Rate

This table presents the results of ordinary least squares regressions in which the left-hand-side variable is the logarithm of annual salary and the right-hand-side variables are as shown. The contribution rate is the employee contribution rate, disregarding contributions out of bonuses and employer contributions. Employee contribution rates less than 4% are recoded to be equal to 4%, and employee contribution rates greater than 18% are recoded to be equal to 18%. All variables are measured as of tenure month twelve for each employee. The sample is composed of the 671 employees whom we observe in the data for at least twelve months. Robust standard errors are in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Indicator for contribution rate equal to 12%	-0.351***	-0.321***	-0.348***	-0.304***	-0.363***	-0.315***	-0.349***	-0.297***
	(0.051)	(0.048)	(0.065)	(0.063)	(0.071)	(0.069)	(0.082)	(0.080)
Contribution rate (percent of pay)	0.042***	0.030***	0.040	0.018	-0.012	-0.018	0.042***	0.026***
	(0.004)	(0.004)	(0.027)	(0.026)	(0.107)	(0.102)	(0.010)	(0.010)
Contribution rate squared (÷100)			0.010	0.052	0.573	0.443		
			0.124	(0.120)	(0.011)	(1.080)		
Contribution rate cubed (÷10,000)					-1.764	-1.227		
					(3.540)	(3.387)		
Contribution rate × indicator for contribution rate >12%							0.001	0.009
							(0.025)	(0.024)
Female		-0.199***		-0.198***		-0.198***		-0.198***
		(0.037)		(0.038)		(0.038)		(0.038)
Married		0.113***		0.115***		0.114***		0.115***
		(0.041)		(0.042)		(0.042)		(0.042)
Age (years)		0.013***		0.013***		0.013***		0.013***
		(0.003)		(0.003)		(0.003)		(0.003)
Fixed effects for first tenure month	No	Yes	No	Yes	No	Yes	No	Yes
R^2	0.150	0.287	0.150	0.287	0.150	0.287	0.150	0.287
Sample size	N = 671	N = 671	N = 671	N = 671	N = 671	N = 671	N = 671	N = 671

Figure 1. A Two-Period Model of the Employee's Contribution Rate Decision

This figure illustrates the structure of employer matching contributions in the retirement savings plan we study. The firm matches employee contributions between 12% and 18% of pay on a one-for-one basis. In this stylized two-period model, income in the present period is one, and income in the future period is zero. The rate of return on contributions is zero. The solid line gives the employee's budget set, and the dotted lines give two possible indifference curves, with their tangency points indicated by circles. The triangle marks the kink in the budget set at the contribution rate of 12%, which is also the default contribution rate. Note that no indifference curve is tangent to the budget set at the default.

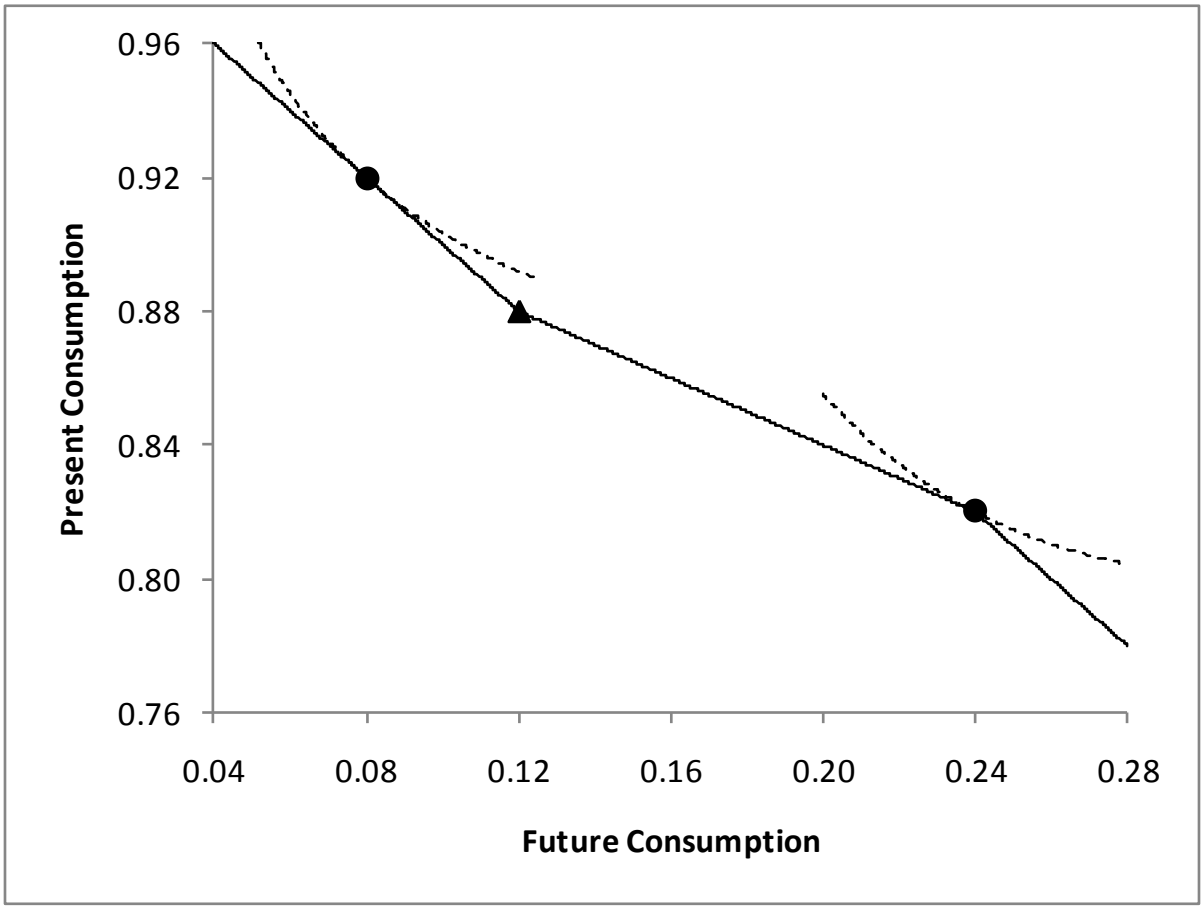

Figure 2. Opt-out from the 12% Default Contribution Rate by Tenure

For a given level of tenure, this figure displays the fraction of employees who have never opted out of the 12% default contribution rate, the fraction who have opted out to a lower contribution rate, the fraction who have opted out of and subsequently returned to the 12% default contribution rate, and the fraction who have opted out to a higher contribution rate. The sample is limited to the 671 employees whom we observe in the data for at least twelve months.

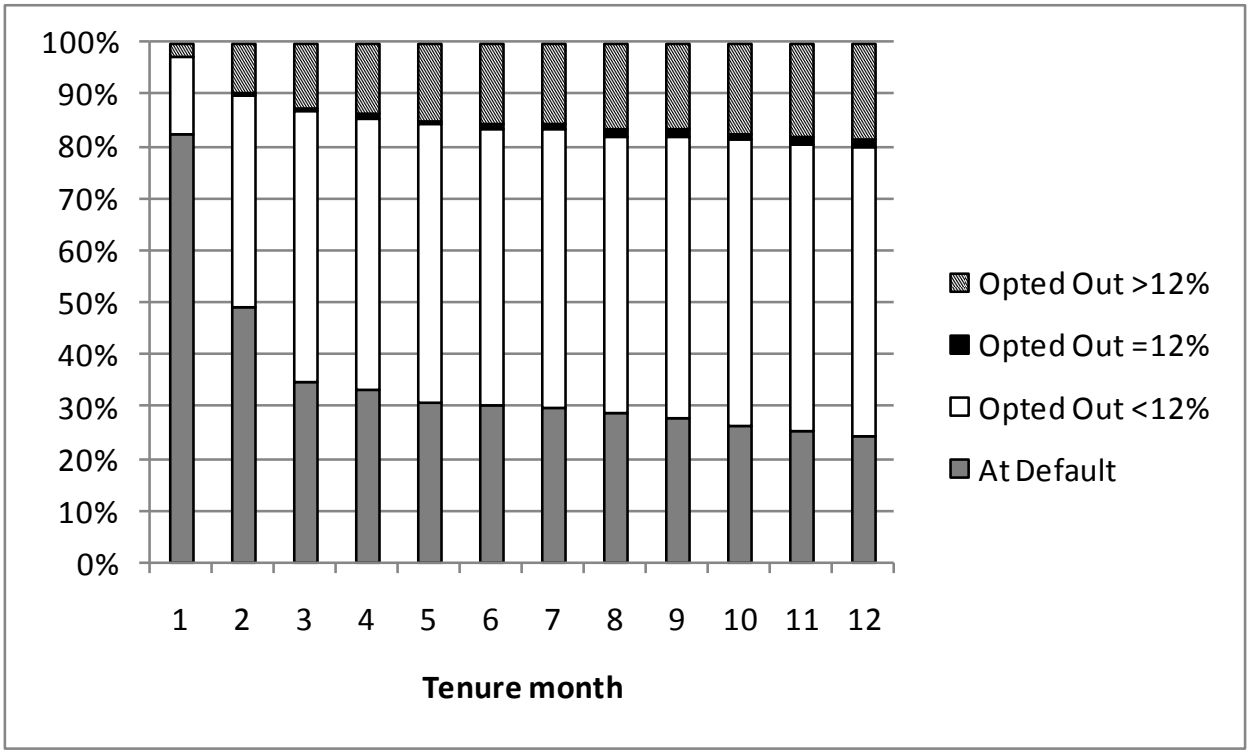

Figure 3. The Distribution of Employee Contribution Rates at Tenure Month Twelve

This figure gives the distribution of employee contribution rates at tenure month twelve in the retirement savings plan we study. Employee contributions out of bonuses and employer contributions are disregarded. The sample is composed of the 671 employees whom we observe in the data for at least twelve months.

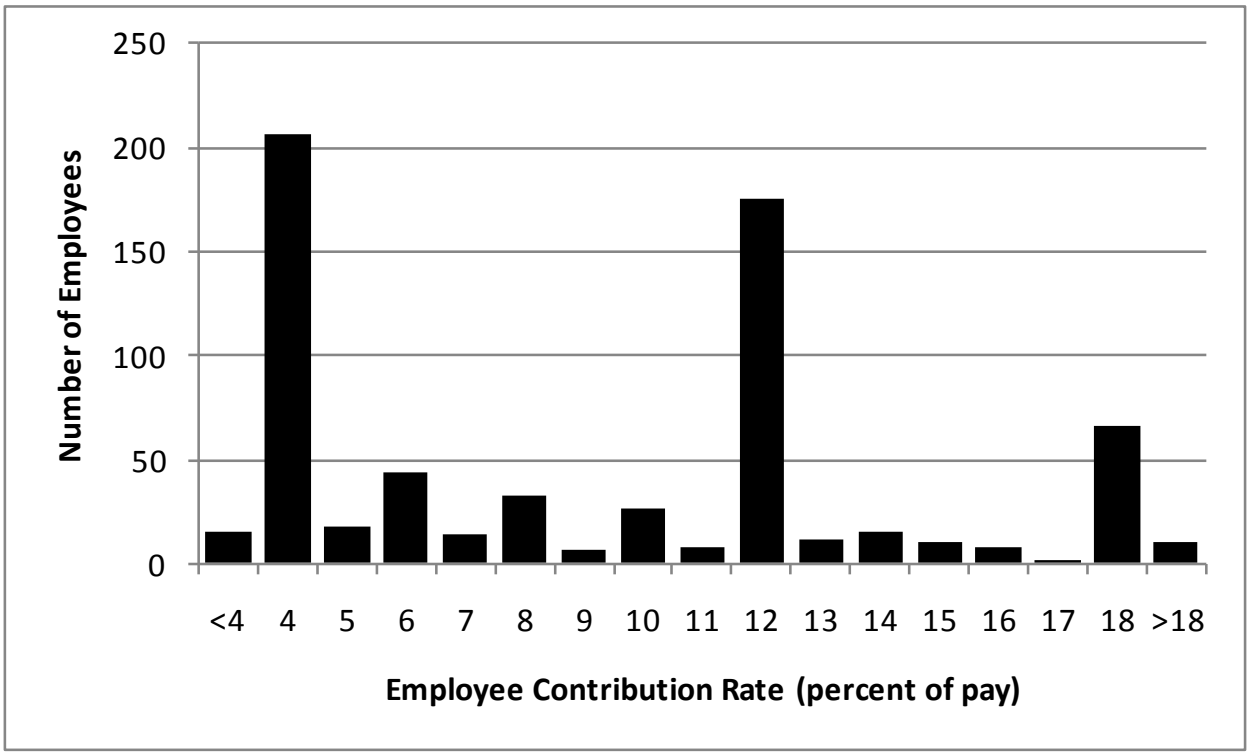

Figure 4. Opt-out from the Default Asset Allocation by Tenure

For a given level of tenure, this figure displays the fraction of employees who have never opted out of the default asset allocation, which is a mix of bonds and equities. The sample is limited to the 671 employees whom we observe in the data for at least twelve months.

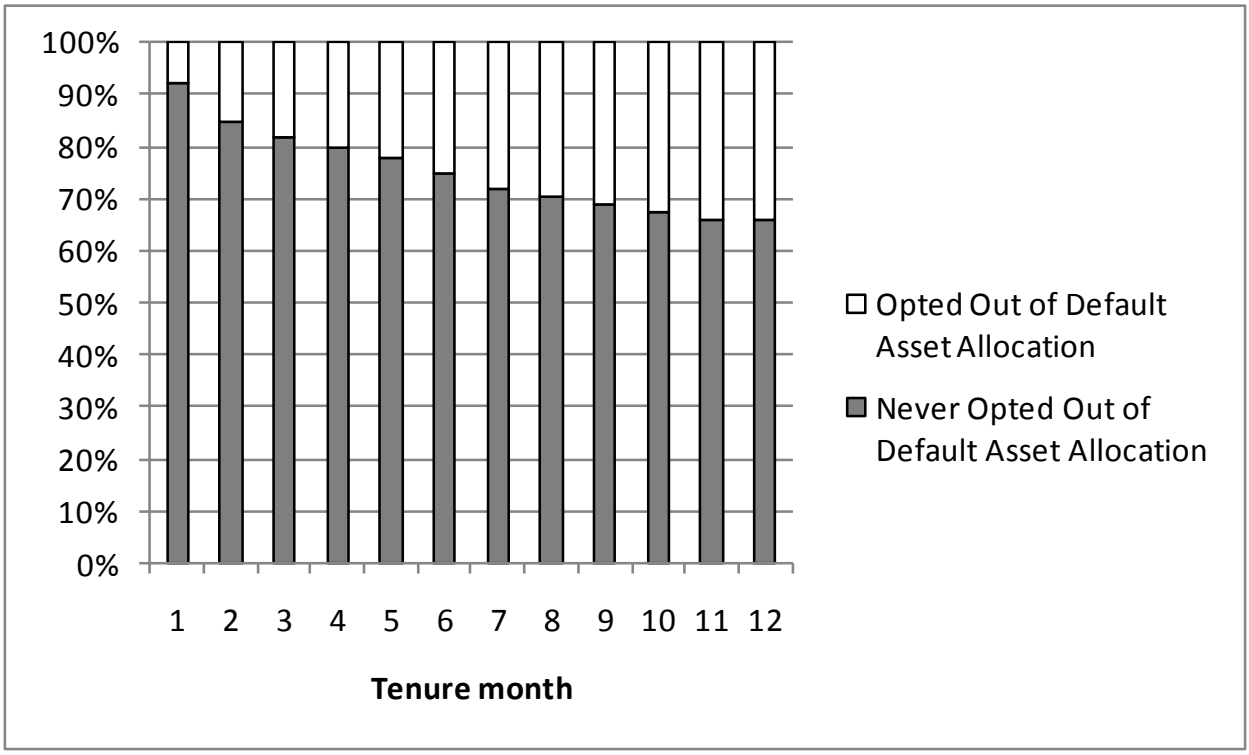

Figure 5. Opt-out from the Default Contribution Rate and Asset Allocation by Tenure

For a given level of tenure, this figure displays the fraction of employees who have opted out of neither the 12% default contribution rate nor the default asset allocation, the fraction who have opted out of the default contribution rate but not the default asset allocation, the fraction who have opted out of the default asset allocation but not the default contribution rate, and the fraction who have opted out of both the default contribution rate and the default asset allocation, which is a mix of bonds and equities. The sample is limited to the 671 employees whom we observe in the data for at least twelve months.

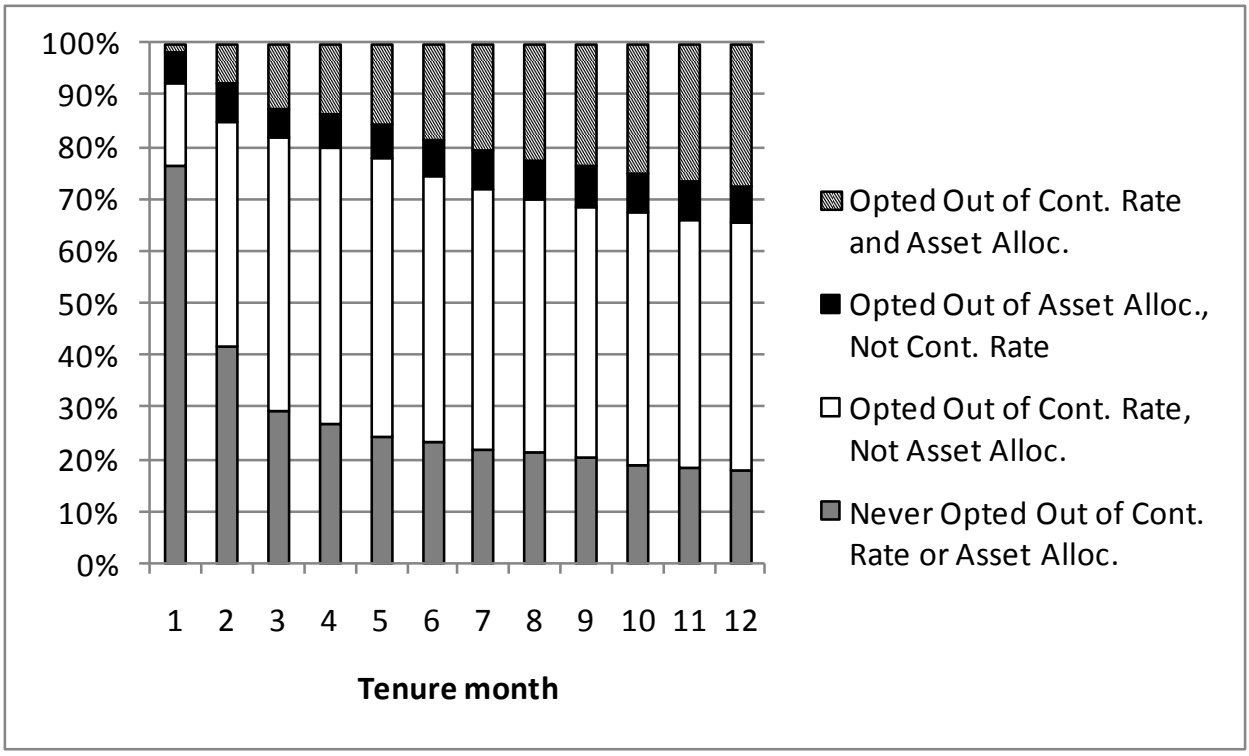

Figure 6. Employee Salaries by Contribution Rate at Tenure Month Twelve

This figure divides employees into groups based on their employee contribution rate at tenure month twelve. Employee contributions out of bonuses and employer contributions are disregarded. Employee contribution rates less than 4% are recoded to be equal to 4%, and employee contribution rates greater than 18% are recoded to be equal to 18%. The boxes display the mean of the logarithm of annual salary for employees in a given group. We also perform an ordinary least squares regression of the logarithm of annual salary on the employee contribution rate, the employee contribution rate squared, and an indicator variable for the employee contribution rate being exactly 12%. The solid line gives the predicted values from this regression, restricting the indicator variable to be zero at all contribution rates. The dotted lines represent two standard errors on either side of the prediction. The sample is limited to the 671 employees whom we observe in the data for at least twelve months.