Balancing Commitment and Liquidity: Empirical Evidence from Mandatory Retirement Saving*

Leora Friedberg
University of Virginia

Adam Leive
University of Virginia

Wenqiang Cai
PriceWaterhouseCoopers LLP

July 2021

Abstract

Recent evidence has questioned the limits of nudges as a means to increase retirement saving. We use administrative data from a large public university to study employee responses to a substantial increase in mandatory retirement saving. We observe full crowd-out of voluntary saving for low-contributing low earners and high-contributing high earners, and otherwise little crowd-out. We interpret crowd-out among low earners as reflecting demand for liquidity and lack of crowd-out as reflecting demand for commitment. We approximate life-cycle welfare consequences of forced saving and find that the benefits to commitment considerably outweigh the costs of reduced liquidity.

JEL Codes: D14, G51, H24, H31, H55

*We thank John Beshears, Taha Choukhmane, Brent Davis, Lee Lockwood, Edgar Olsen, John Pepper, John Sabelhaus, Jason Seligman, and participants of the UVA Applied Microeconomics Workshop, TIAA Fellows Institute Symposium on Behavioral Finance, and 2020 National Tax Association Meeting for their helpful comments. Jiafeng Wu provided excellent research assistance. This project received funding from the TIAA Institute and the Wharton School's Pension Research Council/Boettner Center, and from the UVA Bankard Fund for Political Economy. The content is solely the responsibility of the authors and does not represent official views of the above-named institutions.

1 Introduction

Most employees today with an employer pension have a defined contribution (DC) plan.[1] One of the key features of DC plan design is the considerable autonomy that rests with individuals over how much to contribute, which places major responsibility on them to ensure the adequacy of their retirement income. To nudge people who may be inattentive, uncertain, or dilatory about saving decisions, default enrollment has become the norm (T.RowePrice 2020, Clark and Young 2018). Defaults help employees overcome failures in saving behavior while still permitting choice that allows employees to opt out (Thaler and Sunstein 2008). Extensive research demonstrates the power of defaults in increasing short-term participation in DC pensions (Madrian and Shea 2001, Thaler and Benartzi 2004), but more recent evidence casts doubt on their effectiveness in raising savings rates over the long run (Choukhmane 2019, Laibson 2020). Related theoretical research highlights stronger paternalism mechanisms, especially mandatory contributions, that can raise both long-term saving and welfare under fairly general conditions (Beshears et al. 2020, Moser and Olea de Souza e Silva 2019, Amador et al. 2006). A potential role for mandatory contributions therefore emerges from both recent evidence and theory.[2]

Our understanding of the optimal level of mandatory contributions remains elusive, however. Theoretical models highlight the trade-off between liquidity – valued by people who are unable to borrow against rising future earnings or to finance emergency spending needs – and commitment – valued by people who have trouble adhering to a savings plan. We analyze changes in voluntary DC contributions in response to an increase in mandatory contributions to better understand the demand for liquidity versus commitment. We use administrative data from a public university that instituted an increase in mandatory contributions that was notably greater than the increases studied in other recent settings (Chetty et al. 2014). We interpret the extent to which mandatory contributions raise total DC plan saving by employees, as opposed to simply crowding out their voluntary saving, as revealing information about the demand for liquidity and commitment.[3]

[1] Since the decline of employer-sponsored defined benefit (DB) plans (Friedberg and Owyang 2002), the value of assets in private sector DC plans increased from $74 billion in 1975 to over $5 trillion in 2013 (Saad 2017)

[2] A mandatory savings policy may also serve to rescue society from the "Samaritan's Dilemma" (Buchanan 1975): some people may rationally undersave with the expectation of being helped in old age by an altruistic government (Feldstein 2005).

[3] While we generally refer to a demand for commitment, both a commitment failure (understanding how much to save but not doing so) or an optimization failure (failing to understand how much to save) may explain sub-optimal saving, and mandatory saving can help remedy either. The focus on commitment failures is particular to the literature on defined contribution pensions because of the responsiveness observed to defaults, at least in the short-run.

The increase in mandatory contributions at the university we study is both large and plausibly exogenous. Starting in 2010, newly hired faculty experienced two changes relative to faculty hired earlier: the employer contribution rate fell from 10.4% to 8.9% of salary (a 1.5 percentage point (pp) decline); and, a new mandatory employee contribution of 5% of salary was established (effectively, a 5 pp increase). Consequently, the total mandatory contribution rate rose from 10.4% to 13.9%, a 3.5 pp increase. As predicted by standard economic theory that we review later, the full crowd-out benchmark for optimizing savers is a decline of 3.5 to 5 pp at the high end of the savings distribution (depending in part on whether salaries adjust or not for the change in deferred compensation), and a decline to 0 pp at the low end, as low savers cannot fully crowd out the new mandatory contribution.

Our primary research design compares voluntary DC contributions of employees hired after an increase in mandatory contributions, to employees hired before. This design follows much of the literature on defaults and other retirement savings policies (Madrian and Shea 2001, Carroll et al. 2009, Beshears et al. forthcoming). The policy change followed state legislation focused on funding the state defined benefit (DB) plan, so we view it as exogenous for new employees in the DC plan, rather than, for example, an endogenous response to employee preferences.

We conduct additional analyses and implement an alternative research design to support our interpretation that we are measuring the causal effect of the policy. Using the Survey of Consumer Finances, we demonstrate that there were no observable changes in DC pension plan contribution rates for a similar sample of nationally representative employees. Our findings are robust to different regression specifications and time periods studied. In a placebo test, we fail to detect any change in the choice to invest with either of the two vendors in the DC plan, which is a saving decision that should arguably not be impacted by the change in mandatory contributions but might be if savings preferences changed. Finally, we use academic staff—who were not directly affected by the policy change—as a control group in a difference-in-differences design. Using a control group helps to address concerns that the results may be driven by employer-specific changes in hiring or information provision or by economy-wide changes in retirement savings preferences. We find quantitatively similar results using this alternative research design.

Comparing new employees hired after the policy change to new employees hired before, we find strong evidence of incomplete crowd-out in the aggregate. Voluntary contribution rates overall fell by about 1.5 pp, and we can reject the full crowd-out benchmark of 3.5-5 pp.[4] Participation in the voluntary account fell a little, but not by

[4]If total compensation from the employer remained constant in equilibrium, then an increase in other compensation would accompany the decrease in the employer contribution. We find that inflation-adjusted

enough to explain incomplete crowd-out. Recognizing that deviations from this benchmark may reflect commitment failures, we then consider the full distribution of responses so that we can distinguish between optimizing savers who actively crowd-out the new required contributions and passive savers who do not change their saving.

The distributional analysis allows us to categorize four types of responses to the increase in required contributions that we observe:

1. *Active high savers*: employees saving more than the mandated increase who respond with crowd-out; their welfare is unaffected by the policy.

2. *Active low savers*: employees saving less than the mandated increase who stop making voluntary contributions; they value liquidity and are harmed by the policy.

3. *Passive savers*: employees who display little or no crowd-out and continue making voluntary contributions; they value commitment and are helped by the policy.

4. *Non-savers*: employees who did not make voluntary contributions in the first place; since no actions reveal their type, we consider alternately that they value liquidity and are harmed by the policy or value commitment and are helped.

In order to operationalize this approach, we consider where crowd-out originates along the distribution of voluntary contribution rates, after splitting the sample between high versus low earners. We split the sample because we observe distinctive responses above and below median earnings, which are about $75,000, and further because the value of both liquidity and commitment may differ by income. When we do this, we observe that crowd-out occurs largely at the high end of the savings distribution for high earners (active high savers), and the low end for low earners (active low savers). For the complementary part of these savings distributions, we observe very little responsiveness (passive savers). Overall, we find that active high savers, who optimize with full crowd-out, comprise 17% of our sample; active low savers, who stop participating because, in our interpretation, they demand liquidity, comprise 8%; passive savers, who maintain the same level of voluntary contributions as previously because they demand commitment, comprise 50%; and non-savers, who might demand liquidity or commitment, comprise 25%. Thus, the incomplete crowd-out that we observe for the whole sample is explained by a relatively small share of employees, among both high and low earners, who respond fully.

Given this finding that an increase in mandatory employee contributions is effective at raising total retirement contributions, we evaluate the welfare consequences for those who

salaries rose insignificantly, and we cannot reject a fully-offsetting rise of 1.5 pp. If that occurred, then we would predict full crowd-out of 5 pp.

value liquidity versus commitment. We parameterize a stylized two-period model of risk-averse savers to distinguish the effect of the increase in mandatory contributions on people who face liquidity constraints on the one hand (and are now forced to consume even less while working and even more in retirement) and people who face commitment failures on the other (and otherwise consume too much while working and too little in retirement).[5] Under conservative assumptions, we calculate that the per-employee benefits to commitment for passive savers (at $203 per employee per year) outweigh the costs of liquidity for active low savers (at $170 per employee per year), because some of their reduction in saving is inframarginal and because, even though active low savers have lower earnings on average, their change in saving is also typically smaller. Further, with the share of passive savers exceeding the share of active low savers, the total benefits are at least 1.8 times greater than the total costs and rises by an order of magnitude if non-savers also benefit from the increased commitment.[6]

As noted earlier, our results contribute to research on optimal pension plan design in the presence of commitment or optimization failures. Theoretical and/or quantitative models in Amador et al. (2006), Beshears et al. (2020), Moser and Olea de Souza e Silva (2019), and Bubb and Warren (2020) motivate the use of required or matching contributions.[7] However, we are aware of only two relevant empirical papers, each documenting passive saving behavior in response to mandated contributions. Chetty et al. (2014) study the impact of a 1 pp increase in mandatory contributions in Denmark and find less than a 0.2 pp reduction in pension saving. Notably, we find a similar percentage of active responders, of around 17% in our case, who exhibit full crowd-out, even though the mandatory increase in our setting is five times larger.[8] In related work to ours, Card and Ransom (2011) study how the mix of required contributions impacts voluntary contributions among faculty using variation across 77 universities between 1986 and 1996. They estimate crowd-out rates of between 60 and

[5] We assume, conservatively, that in the first case the now-liquidity constrained employees were just barely unconstrained and therefore saving optimally ex ante, and that in the second case the now-committed employees are just barely surmounting their commitment problem and therefore saving optimally ex post. Otherwise, the changes in welfare, negative for the first group and positive for the second, will be greater.

[6] This assumes a coefficient of relative risk aversion of 3, while treating the baseline non-participants as liquidity-constrained rather than non-optimizing. The ratio of benefits to costs falls modestly if risk aversion is higher.

[7] Bubb and Warren (2020) differs from the others in studying employer rather than social-planner motives for paternalistic pension design when some individuals experience commitment failures. Their equilibrium search model suggests further limitations of nudges, as employers offer default enrollment at a contribution rate below the employer matching rate. Naive employees are attracted by the high matching rate, accepting a lower wage in exchange but, in the end, do not take advantage of it, while sophisticated employees benefit from it and obtain higher compensation as a result.

[8] This provides suggestive evidence against rational inattention as an explanation, as a larger change in mandatory contributions should elicit a greater response, assuming that the underlying distribution of optimal savings rates is similar.

80 percent of the theoretical benchmark in response to mandated employee contributions, but rates only half as large in response to employer contributions. We cannot separately distinguish the response to the mandated employee versus employer contribution in our setting, yet our results give further insight into incomplete crowd-out by exploiting variation within a single university over time and during a much more recent period. Our substantially lower estimated crowd-out rate, of about 30 percent, may reflect an increasingly complicated choice environment that employees face across their benefits plans.

The rest of this paper is organized as follows. Section 2 describes the institutional details of the large public university recorded in the data. Section 3 introduces a theoretical model that predicts how the average contribution rate to voluntary plans changes when the total contribution rate to the mandatory plan increases. Sections 4 to 6 describe, respectively, the data, empirical strategy, and the main results of this paper. Section 7 analyzes the welfare effects of the policy and Section 8 concludes.

2 Institutional Details

The large public university that we study offers faculty a complicated set of retirement plan choices. The flow chart in Figure 1 summarizes the sequence of retirement plan choices and available options. First, faculty face a one-time irrevocable choice at the outset of employment between the DB plan run by the state and the 401(a) DC plan with the mandatory contributions that we described earlier.[9] As shown in Section 4, about 80% choose the DC plan, a fraction that rose a little but not significantly, in spite of the new mandatory employee contribution. Second, faculty must choose, for their mandatory contributions, among two vendors and numerous funds for each vendor.[10] Third, faculty can choose voluntary contributions with a very limited match rate, of 50% for up to $40 per month, which go into the university-run 403(b) plan. Fourth, faculty can choose additional voluntary contributions, which can be directed to the same 403(b) plan and to a state-run 457 plan, and, near the end of our sample period, to a university-run Roth 403(b) plan and a state-run Roth 457 plan.[11] Fifth, faculty must choose among two vendors for

[9]We will refer to the 401(a) plan as the "university plan with required contributions" or the "mandatory DC plan", though it is only mandatory after the DB-DC choice; this terminology helps distinguish it from the voluntary DC plans that are our main focus. The state DB plan became a hybrid DB-DC plan in 2014, but that occurs in the last year of our sample and did not noticeably change the DB plan choice.

[10]Employees can split contributions among the two vendors and among all the funds offered by each vendor.

[11]The 403(b) and Roth 403(b) options are jointly subject to IRS contribution limits, just like 401(k) plans are. The 457 and Roth 457 options are jointly subject to additional IRS contribution limits, meaning that faculty are able to contribute twice as much to retirement plans as employees who work for most other types of employers are.

403(b) contributions, while the 457 plan has a single (different) vendor, and among numerous funds for each vendor, with fund menus differing for the 401(a), 403(b), and 457 plans. The focus of our paper is on the sum of voluntary contributions made in the third and fourth steps.

The voluntary DC plan likely represents the most attractive account to save for retirement due to its tax preferences, low expenses, and ease of taking loans. The advantages of the 403(b) and 457 relative to IRAs or taxable accounts is important since we only observe employer accounts. We therefore assume that the overall savings rate responds to the change in required contributions, but the choice of where to save does not.

Conditional on choosing the university DC plan in the first step above, the total required contribution rate in the plan depends on the employees' hiring date, as detailed in Table 1. For employees hired before July 1, 2010, the university contributed 10.4% of their monthly pay to the DC plan, and employees were not required to contribute any money to the plan. State legislation enacted in 2010 implemented modest measures to enhance funding of the state DB plan, while also mandating a 5% employee contribution for those few state employees (mostly faculty) who could opt for a DC plan, to match the employee contribution rate already required for the DB plan.[12] In response, the university reduced its employer contribution rate to 8.9% for employees hired after July 1, 2010, resulting in a total contribution rate of 13.9%. This change in mandatory contribution rates can thus be viewed as an exogenous change for new faculty, rather than an endogenous response to new faculty preferences or an element of a broader set of changes in university policy.

We exploit the exogenous shift in the mandatory contribution rates to investigate whether newly hired faculty who choose the DC plan alter their voluntary contribution rates in response. A few additional comments about the voluntary contribution setting are in order. As noted above, a small employer match is available, of up to $20 per month for voluntary contributions of $40. Beginning in 2009, the university began to auto-enroll employees in the 403(b) plan with the default vendor at a contribution of $40, so that they would automatically benefit from the match.[13] Our empirical approach will account for this change during our sample period from a default of $0 to $40 in voluntary contributions by treating anyone who contributes at a very low rate, whether before or after the default

[12] The changes in the DB plan involved raising the eligibility age for benefits, slightly reducing service credits, and reducing the cost-of-living adjustment once benefits are being received. These changes reduce effective DB pension wealth. For faculty, it is not clear which changes would weigh more in the initial choice between the DB and DC plan, and, as we note later, we observe a small and insignificant response in this choice. For staff, whom we use as a control group in an alternate specification to our event study, reduced DB pension wealth should lead to more voluntary contributions through a crowd-in effect or less contributions through a reduced-wealth effect; and, in any case, we observe almost no change in behavior.

[13] If the employee actively chose the second vendor, there was no default contribution rate.

was enacted, as not contributing; this recognizes the absence of an active choice on their part. Later, in 2013, the Roth options mentioned above became available. Because the Roth option does not change the maximum contribution amounts that may be made, we treat voluntary contributions to either the TDA or Roth accounts as equivalent, after adjusting for an imputed employee marginal tax rate.

The next section presents a simple theoretical model that predicts how employees respond to changes in mandatory contribution rates, and the two sections after introduce our data and empirical strategy to test whether employees respond as standard theory predicts.

3 Model

This section introduces a stylized theoretical model of optimizing savers to explain how changes in required employer and employee contributions should affect employees' participation and contribution decisions in their voluntary plans. Understanding conditions under which mandatory contributions are predicted to crowd out voluntary contributions lets us establish a full crowd-out benchmark that we use to evaluate whether observed crowd-out responses are consistent with optimizing behavior.[14] We will also use this model later to evaluate welfare costs in case employees experience liquidity constraints (because they are forced by the mandatory contribution to consume too little in the first period) or commitment failures (and end up consuming too much in the first period).

We highlight three predictions of the benchmark model for voluntary contributions, along with welfare effects. One is a crowd-out effect on voluntary contributions, from the new required employee contribution. The second is a compensation effect, from the effective change in total employer compensation resulting from the reduction in the required employer compensation, to the extent that this is not compensated by an increase in salary. The third is a participation effect, which constrains the crowd-out effect for employees who would otherwise contribute amounts that are too small to generate full crowd-out. The compensation effect reflects a welfare loss for all employees, while the participation effect reflects an additional welfare loss because some employees are forced to oversave because of voluntary contributions.

[14]We refer to the optimizing saver predictions as the full crowd-out benchmark to distinguish it from the partial crowd-out alternative, in which employees respond by less than what is predicted. Nevertheless, as we show next, some optimizing agents respond by less than others because they have less savings to crowd out; in their case, full crowd-out means that their voluntary contributions drop to zero.

3.1 Baseline Model with Universal Saving

We begin with a model in which all employees have strong savings motives, so their optimal savings exceed the mandatory contribution. An employee lives for two periods. She works during the first period and earns income Y and retires during the second period. An employee maximizes her lifetime utility subject to the following constraint: since no borrowing is allowed, consumption in each period cannot exceed resources available in that period. The only way to finance consumption in the second period is by spending down retirement saving. In the baseline model, we assume that an employee's saving in the absence of any required contributions exceeds the required contribution amounts.

Employees must contribute M_e to the mandatory retirement savings account, and the firm must contribute M_f. Employees then have a choice between contributing an additional non-negative amount M_v to a voluntary retirement savings account.[15] We assume that both mandatory and voluntary retirement saving accounts have the same rate of return r. We further assume no change in income Y, so that effective compensation declines.[16]

Before the policy change, the firm is required to make a contribution to the retirement account, but the employee's required contribution $M_e = 0$. Therefore, employee i hired before the policy change solves the following maximization problem:

$$\max_{c_1} u_i(c_1^i) + \delta u_i(c_2^i)$$
$$\text{s.t.} \quad c_1^i = Y - M_v^i$$
$$c_2^i = (1+r)(M_v^i + M_f)$$

where $u_i(\cdot)$ is a concave function and δ is the discount factor. First-period consumption c_1^i equals first-period earnings Y minus the voluntary contribution M_v^i that employee i chooses. Meanwhile, the constraint against borrowing means that the employee cannot access the employer's mandatory contribution until retirement. Second-period consumption c_2^i equals the accumulated balance of the retirement account, which consists of the employer contribution M_f and the voluntary employee contribution M_v^i, along with interest earned on those contributions.

Now, consider employee j who is hired after the policy change, which reduces the

[15]We ignore income taxes and the associated tax deferral available through the retirement plan; legal limits on voluntary DC contributions; and the very small default contribution of $40 and associated matching contribution of $20 per month. These do not alter substantive predictions of the model.

[16]If instead, salaries increase by 1.5%, then the compensation effect outlined below would not operate, leading to a bigger predicted crowd-out effect. In fact, we observe a modest increase in imputed salary in our sample and cannot reject either 0 or 1.5%, but since our salary data is reported in ranges, we are not confident in claiming that no compensation effect arises.

required employer contribution and establishes a required employee contribution, leading to a higher required total. She maximizes the same objective function as employee i but faces modified budget constraints:

$$
\begin{aligned}
\max_{c_1} \; & u_j(c_1^j) + \delta u_j(c_2^j) \\
\text{s.t.} \quad & c_1^j = Y - M_e^j - M_v^j \\
& c_2^j = (1+r)(M_v^j + M_f + M_e^j)
\end{aligned}
$$

The new required employee contribution M_e now appears negatively in the first-period budget constraint and positively in the second-period budget constraint.

Figure 2 is a graphical representation of the baseline model using actual numbers from the policy change. BL^i represents the budget line before the policy change. Applying the budget constraints for employee i, the two end points for BL^i are $c_1 = Y, c_2 = (1+r)0.104Y$ and $c_1 = 0, c_2 = (1+r)1.104Y$, because the firm's contribution before the policy change is 10.4% of monthly pay. Similarly, BL^j is the budget line after the change. The two endpoints are $c_1 = 0.95Y, c_2 = (1+r)0.139Y$ and $c_1 = 0, c_2 = (1+r)1.089Y$. This makes it clear that the policy change reduces employee j's resources relative to employee i because of the reduction in the employer contribution of 1.5 pp – and not because of the new required employee contribution of 5%.

We assume that preferences of employees hired before versus after the policy change do not differ systematically and that consumption in both periods is a normal good. For a particular set of preferences, point A is an example of the optimal consumption bundle for employee i, hired before the policy change, and point B is an example for employee j, hired after the change. The distance between Y and c_1^i is M_v^i, the voluntary contribution amount for employees hired before the policy change. Similarly, the distance between $0.95Y$ and c_1^j is M_v^j, the contribution amount to the voluntary plan for employees hired after the policy change.

The change in voluntary plan contributions is then:

$$
\begin{aligned}
M_v^j - M_v^i &= 0.95Y - c_1^j - Y + c_1^i \\
&= -0.05Y + c_1^i - c_1^j
\end{aligned}
\tag{1}
$$

Equation (1) demonstrates two effects on voluntary contributions. The crowd-out effect is reflected in the first term, $-0.05Y$, as the new required employee contribution makes it unnecessary to save the same amount in voluntary form. The compensation effect is reflected in the second term, which is positive (since c_1 shrinks) and therefore undoes part of the crowd-out effect. Because consumption in each period is a normal good, employees want to reduce consumption a little in both periods rather than experiencing a larger drop in period 2, and they do this by saving a little more than they would otherwise in period 1.

A necessary condition for the crowd-out effect to operate is that the employee chose to save at least 3.5% salary before the policy change. This is represented by point A, showing the optimal location for employee i, lying to the left of point D along BL^i. If so, then $0 \leq c_1^i - c_1^j \leq 0.015Y$ since BL^j is a parallel shift inwards of BL^i.

Applying this inequality to equation (1) yields:

$$\begin{aligned} -0.05Y \leq M_v^j - M_v^i &\leq -0.05Y + 0.015Y \\ \iff -0.05 \leq \frac{M_v^j - M_v^i}{Y} &\leq -0.035 \end{aligned} \tag{2}$$

Therefore, the model predicts that the reduction in the voluntary contribution rate lies between 3.5 and 5 pp. This bounds the compensation effect between 0 and 1.5 pp.

3.2 Incorporating Non-Participation

The predictions in our baseline model pertain to the case of relatively high voluntary contribution rates. If voluntary contributions for employees before the policy change are small or zero, however, that limits the crowd-out effect, replacing it instead with a participation effect.

In Figure 2, this case is illustrated by employee with indifference curve U_L^i, who locates at point D before the policy change. After the policy change, all employees with low enough counterfactual contribution rates will bunch at E, which is shown by U_L^j locating at the notch in BL^j. In this case, $c_1^i - c_1^j \leq Y - 0.95Y = 0.05Y$. Following the same analysis as earlier,

$$\begin{aligned} M_v^j - M_v^i &= -0.05Y + c_1^i - c_1^j \\ \implies -0.05Y \leq M_v^j - M_v^i &\leq -0.05Y + 0.05Y \\ \implies -0.05 \leq \frac{M_v^j - M_v^i}{Y} &\leq 0 \end{aligned} \tag{3}$$

For those who cannot fully offset the crowd-out effect, contributions fall by an amount between 0 and 5 pp that is smaller than above. This participation effect represents the truncated outcome of the crowd-out effect, since employees who cannot reduce their contributions below zero instead stop contributing entirely to the voluntary retirement plan.

Among these employees, mandating employee contributions results in a welfare loss that exceeds what arises due to the compensation effect. This is shown by the distance between U_L^j and U_L^i. The participation effect results in forgone consumption in period 1 of up to 5% of salary. By contrast, the utility losses for employees with high saving rates are much smaller: those saving above 5% salary can fully offset the mandatory increase through reduced voluntary contributions, and only experience the small utility loss due to the compensation effect.

3.3 Discussion

Table 2 summarizes theoretical predictions for the full crowd-out benchmark with optimizing savers. For employees who would save enough counterfactually, the average voluntary contribution rate will decrease by 3.5 to 5 percentage points, as any compensation reduction partially offsets the crowd-out effect. For employees who would contribute little or nothing voluntarily, the average voluntary rate is predicted to decrease by smaller amounts (because some cannot reduce contributions by as much as they would like), while the average voluntary participation rate is also predicted to decrease. Such employees are forced to oversave by mandatory contributions, which reduces their welfare. Since the participation responses are predicted to occur among employees who would optimally not save much otherwise, we will investigate the distribution of responses and not just the mean change in voluntary contributions. Moreover, since lower earners are more likely to face binding liquidity constraints, we will split our sample by earnings when investigating the distributional responses.[17]

As we discuss later, failures to optimize saving behavior may change both the predicted responses to mandatory contributions and their welfare consequences (Amador et al. 2006, Beshears et al. 2020, Moser and Olea de Souza e Silva 2019). While the model above assumes perfect rationality and no frictions in decision-making, considerable evidence related to pensions suggests that individuals do not fully optimize their saving behavior (Madrian and Shea 2001, Thaler and Benartzi 2004, Thaler and Sunstein 2008). Some of these failures may be explained if people have trouble committing to a long-term savings plan. Meanwhile, related evidence in saving and other settings suggests not only commitment failures but also optimization failures, arising either because people

[17] Yet, those with high incomes may also face liquidity constraints, as documented by Kaplan et al. (2014).

misunderstand the planning environment (Friedberg and Leive 2021) or cannot solve complicated dynamic programming problems under uncertainty.[18] In case of either commitment or optimization failures, mandatory contributions may raise welfare by helping to increase total savings rates. We thus use observed responses to the increase in mandatory contributions to distinguish among these possibilities.

4 Data

To study responses to mandatory contributions and what they reveal about the demand for liquidity versus commitment, we use administrative data from an employer defined contribution plan.

4.1 Data Description

We construct a novel panel data set using a large university's administrative records. The administrative data contain monthly retirement plan information, semiannual demographics information, and annual earnings collapsed into bins in order to eliminate the possibility that an individual can be identified. Retirement plan information consists of employee and employer contribution rates (as a percent of earnings) to all available retirement savings plans each year.[19] Demographic information consists of employee gender, age collapsed into bins (again, to maintain data confidentiality), marital status (which is incompletely collected), hiring year, and category of employment (faculty versus staff). To control for macroeconomic conditions, we merge in monthly levels of the S&P 500 stock index, inflated to January 2018.

We use an event study approach to compare voluntary contribution rates of new faculty hired in the five years before versus five years after July 1, 2010, when contribution rates were changed. The panel data contain monthly records of approximately 2,800 new faculty hired between 2005-2014. We focus only on contributions in the months of September through May, when faculty receive their full-time pay. We exclude a relatively small number of unusual observations.[20]

[18] We do not consider a formal model of such failures, and it is possible that some individuals, especially those who may be naive about their commitment or optimization failures, might persist in their failures by crowding out mandatory contributions.

[19] While our salary data is collapsed into bins, the contribution rates as a percent of salary are exact numbers.

[20] Starting with records on about 3,000 newly hired faculty, we exclude observations with arrears payments in any retirement savings plan in any year because arrears (when the employee should have contributed to the plan but did not) may reflect salary interruptions that could alter contribution decisions. We also exclude observations with annual income less than $10,000 because these observations probably represent faculty who worked at the university for a very short period of time.

We consider contribution rates of new faculty nine months after being hired, when participants reach a steady contribution rate.[21] In some of our analysis we split the sample by median earnings, resulting in a little over 1,400 new faculty in each group. We choose to focus on participation and contribution rates in the first year since rates of separation after one year are similar across cohorts hired before the change versus hired after. Separation rates beyond two years, however, differ somewhat between these groups, and this leads to some compositional differences in the sample of employees who remain with the University. Nonetheless, we consider outcomes in the first three years in robustness tests and find quite similar overall patterns.

4.2 Descriptive Statistics

The summary statistics in Table 3 use data from the end of the employee's first year to show demographic and employment characteristics for faculty hired before and after the policy change; the last two columns report the mean difference between these two groups of faculty, and the corresponding p-value from the t-test that the means are equal. Appendix Figure A.1 further shows trends in key demographic characteristics over time, which makes it clear that new faculty characteristics did not undergo any meaningful changes at the same time as the policy change.

The demographic characteristics we observe are age (in bands), gender, and marital status, all of which are known to affect the propensity to save for retirement. Demographics across the two groups are generally balanced. The only significant differences that emerge are small and occur for age and earnings, variables we observe inexactly. Because age is grouped into bands (generally of 5 years) in order to protect the identity of individuals in our sample, we impute an exact age using the mid-point of the age bands. The average (imputed) age in our sample dropped slightly from 39.7 before 2010 to 38.9 after. 43% of new faculty in the sample are female. Marital status is reported on an incomplete basis, because the university does not make use of information that it collects at the outset of employment for any personnel purpose, and many employees do not report or update their marital status. Roughly 44% of new faculty report being married and 25% report being single, while the remaining 30% do not report their marital status.

The key employment characteristics that we observe for faculty are full-time status and earnings. The share who are full-time stays the same, at 91%, before and after 2010. With earnings reported in bands, we use the midpoint of each band in calculating the means in

[21] Appendix Figures A.5a, A.5b, and A.5c make it clear that participation and contribution rates remain steady in the second and third year after hire. Madrian and Shea (2001), for example, show that retirement plan participation rates increase over time, apparently as employees overcome initial procrastination.

Table 3.[22] In real terms, average (imputed) salary rose from \$84,373 before 2010 to \$87,473 after, a modest increase that is marginally significant at the 10 percent level. Because we observe salary imprecisely, we cannot say with much confidence whether or not it rose in possible compensation for the decline in the employer contribution. To the extent that it did, the full crowd-out benchmark would be a 5% decline in voluntary contributions, rather than a 3.5% decline – both of which our evidence rejects. In some of our later analysis, we split the sample by median annual salary, which corresponds to approximately \$75,000. Table A.3 splits descriptive statistics by faculty with salaries above or below this amount. On average, higher-salaried new hires are more likely to be men and are about 4 years older than lower-salaried new hires. Few characteristics change significantly or substantially over time within these two groups.

Lastly, the stock market, as measured by the S&P 500, rose substantially during this period. We include this as a control in our regressions, and in Section 6 we exclude the period around the Great Recession as a robustness check and find similar patterns.

4.3 Retirement Plan Choices

Table 3 also shows raw statistics on retirement plan choices before and after 2010. The aggregate patterns that we find in our econometric analysis in Section 6 are evident in the raw data.

Recall from Section 2 and Figure 1 that faculty first face a one-time irrevocable choice between the DB plan run by the state and the university DC plan with mandatory contributions on which we focus. Table 3 shows that the share choosing the DC plan, which is about 20% throughout, rose a little after 2010, by 1.6 pp. While not significant, this increase may reflect the decline in effective DB benefits in the state plan which occurred at the same time, offsetting the increase in the mandatory contribution in the DC plan.

Next, we compare participation in the voluntary retirement plans (summing together both the 403(b) and 457 options) before and after 2010. These statistics adjust for two other changes in retirement plan structure during our time period. In 2009, auto-enrollment into the 403(b) plan began with a very small contribution for new employees; the default monthly contribution was set at \$40 in order to provide employees with the maximum employer match of \$20. To deal with this change in the plan environment, we treat all participants with very small contribution levels that are consistent with the cash match as not making voluntary contributions; this includes employees making an active choice before 2009 involving a small contribution that was matched. We do this so as not to distort our analysis in response to

[22] In particular, salary is reported in \$5,000 bands below \$100,000 annual salary, and then \$20,000 bands to \$200,000, with top-coding above this level.

an increased frequency of very small default contributions; we do not anticipate that anyone contributing just enough to get the cash match will respond to the crowd-out effect of the large shifts in required contributions that are our focus.[23] In 2013, Roth versions of the voluntary retirement plans become available, and we include them (adjusted for their post-tax status using an imputed marginal tax rate) along with traditional tax-deferred accounts as voluntary plan options.[24] Appendix B describes the imputation procedure that we used to construct the marginal tax rate.

Among faculty who chose the university DC plan with mandatory contributions, Table 3 shows the share choosing to contribute to any of the voluntary DC plans at the end of their first year, before and after 2010, and Figure 3a shows the same over the first nine months (the first year excluding summer months) after hire. By the end of the first year, participation rates are slightly lower for those hired after 2010 (shown with dark bars), compared to those hired before (shown with light bars). Figure 3b shows the voluntary contribution rate over the same period among those making positive contributions.[25] Here, new faculty after 2010 chose voluntary contribution rates that were about 1 pp smaller by the end of the first year. Figure 3c plots aggregate voluntary contributions averaged over those who participate and those who do not, which is of ultimate interest when evaluating the adequacy of retirement saving. In combining the participation and the contribution margins, we observe close to a 1 pp decline in contributions for faculty hired after versus faculty hired before the policy change by the end of the first year. These patterns indicate an aggregate crowd-out effect that appears quite small.

While we conduct regression analysis on the means of the contribution variables to mirror the raw statistics above, we also analyze the distribution of contributions, given the theoretical predictions that crowd-out responses should differ for employees who want to save a lot versus a little. In this distributional analysis, we further split the sample by earnings, since liquidity constraints may bind more for lower earners. In Table A.3, which splits descriptive statistics by faculty with salaries above or below $75,000, on average, higher-

[23]Since we do not observe dollar amounts of contributions but rather contributions as a percent of salary, we use the upper bound of the salary bin of each employee to construct the percentage contribution consistent with the cash match, and set to zero any contribution rates that are at or below that amount. For example, for someone in the [$50k, $55k) salary bin, we label 403(b) contributions below 0.8% as being consistent with the cash match and record this contribution as zero. Our results are robust to instead using the lower bound of salary. Since the salary bins are relatively narrow and the cash match amounts are small, the estimates are not sensitive to this definition. We choose the upper bound to be conservative in replacing small contributions with zeros.

[24]While some participants began to choose the Roth rather than tax-deferred options, total participation and contribution rates trended smoothly before and after this date, so it did not seem to alter contribution amounts.

[25]We winsorize voluntary contributions at the 99th percentile (equal to contributing about 50 percent of salary) to remove the influence of outliers on the mean.

salaried faculty are more likely to make voluntary contributions, and save more conditional on contributing. Therefore, we would expect differential crowd-out if all employees make optimal savings decisions, and different valuations of liquidity versus commitment otherwise.

5 Empirical Strategy

In order to evaluate the trade-off between liquidity and commitment, we analyze crowd-out responses to the increase in required contributions that occurred at the university that we study. We begin by estimating regression specifications for the full sample to test for crowd-out in the aggregate. Rejecting that, we analyze the distribution of voluntary contributions to pinpoint crowd-out on the intensive and extensive margins, and we further distinguish among lower and higher earners, who are more or less likely to experience liquidity constraints. Comparing these estimates to the full crowd-out benchmark that we derived earlier for optimizing savers allows us to determine what fraction of employees value liquidity versus commitment.

5.1 Identification

The ideal experiment to estimate employee responses to a mandatory increase in contributions would be to randomly require some employees to make mandatory contributions while exempting others. The difference between the two groups would yield the overall effect of the policy. Our setting differs from this ideal setup in that all the new faculty whom we observe after 2010 experienced the change in contribution rates. Lacking a control group of faculty, we compare participation and contribution rates among new faculty hired before versus after July 1, 2010, similar to the methods used in much of the literature on defaults and other retirement savings policies (Madrian and Shea (2001), Carroll et al. (2009), Beshears et al. (forthcoming)).

The key identification assumption necessary for this approach is that new faculty hired after the policy change do not differ systematically in their saving preferences compared to new hired faculty hired before. We offer evidence that helps rule out several concerns with this assumption. One concern is that the willingness of faculty to accept job offers at the university we study changed as a result of the shift in required contributions. However, as we showed earlier in Table 3 and Appendix Figure A.1, faculty changed little in their observable employment or demographic characteristics, and this reduces concerns that changes in hiring practices or in the selection of faculty accepting job offers changed unobservably at the same time.

A second concern is that other university policies changed concurrently. No other changes in the university DC plan took place, and we have not found evidence of changes in compensation or hiring policies. The state DB plan, which faculty can elect at the outset of employment in a one-time choice, became less generous in 2010, which counteracts the relative compensation decline resulting from the reduced employer contribution in the DC plan. Since overall faculty participation in the DB plan remains low and changes by a small amount, as noted in the previous section, we do not view this policy change as a threat to identification.

A final concern is that other changes in the broader economy occurred at the same time. The July 1 implementation of the policy is standard for state law changes that were enacted months earlier, following weeks of debate, while the underfunding of the state DB pension had been in the news for almost two years, since the financial crash of 2008. Therefore, it is unlikely that other events associated with the financial crisis or with underfunding coincided exactly with the new policy. There were also not major changes to other features of the mandatory and voluntary plans, such as provisions for rollovers, loans, and withdrawals, during the study period.

In order to validate that the timing of our policy change is unrelated to underlying trends in savings behavior or in the broader economy, we examine nationally representative data on retirement plan behavior. We use the Survey of Consumer Finances, which is administered every three years and is typically viewed as the most reliable individual-level data set available on financial behavior. We extracted a sample with similar education and earnings and found that nationwide participation and contribution rates in voluntary DC plans varied little over our sample period. This evidence is described in more detail in Appendix Table A.1.

As supplementary analysis, we treat academic staff as a control group when examining voluntary contributions, to difference out any cohort-specific differences before and after the policy change. Unlike faculty, staff were enrolled in the DB plan and could not choose the mandatory DC plan instead, but they face the same voluntary DC options as before. We do not use this as our main specification because staff were affected by simultaneous changes to the DB plan that modestly reduced DB pension wealth, which might lead to more voluntary contributions through a crowd-in effect or less contributions through a reduced-wealth effect.[26] Using staff as a control group may, nevertheless, help address concerns that the results may be driven by other factors, such as employer-specific changes in hiring or information provision or by economy-wide changes in retirement savings preferences.

[26] As we noted earlier, these changes affected the benefit eligibility age, service credits, and cost-of-living adjustments, but notably did not alter vesting, which may be more salient to new employees.

5.2 Full-Sample Analysis

As noted earlier, we begin by analyzing crowd-out in our full sample. This indicates whether we can reject the full crowd-out benchmark, which predicts reductions in contribution rates of between 3.5-5 percentage points for those contributing enough to fully crowd out the increased required contributions, and a reduction to zero for the rest. Our specification for this analysis accounts for the fact that many participants do not choose their contribution rates right away (Madrian and Shea 2001), so we consider contribution rates over the first year of employment.

We estimate equations of the following form using OLS:

$$\begin{aligned} y_{it} = \alpha + \beta_1 Female_i + \beta_2 Single_{it} + \beta_3 Married_{it} + \beta_4 Fulltime_{it} + \beta_5 Stock_t \\ + \sum_a \beta_a \mathbb{1}[Age_i = a] + \sum_j \beta_j \mathbb{1}[Income_i = j] \\ + \beta_p Post_i + \sum_k \beta_k \mathbb{1}[Tenure = k] + \sum_t \beta_t \mathbb{1}[Tenure = t] \times Post_i + u_{it} \end{aligned} \quad (4)$$

The key explanatory variables in (4) are $\mathbb{1}[Tenure = t] \times Post_i$, which estimate the contribution behavior in each month t of *Tenure* for faculty members hired on or after July 1, 2010, when *Post*=1, compared to faculty members hired before. The outcomes y that we consider are the voluntary contribution rate, as a percent of salary, and the binary decision of whether to make voluntary contributions, for those faculty who chose at the beginning of their tenure to participate in the university DC plan with its mandatory contributions. In addition, we run a preliminary regression for the full sample with the outcome y defined as the one-time irrevocable binary choice to participate in the state DB plan rather than the university DC plan, in this case without the indicators for months of tenure.

In all the regressions, we control for demographic and employment characteristics that might influence retirement contributions. The demographic characteristics, as we discussed in the previous section, are *Female* for gender, *Married* and *Single* for marital status (where the omitted category is unknown marital status), and *Age* for the age ranges that we observe for our sample. The employment characteristics are *Income*, annual income measured in thousands of 2018 dollars and imputed as the mid-point of the income bands that we observe for our sample; and *Fulltime*, an indicator that takes a value of one if the faculty is hired as a full-time employee. Lastly, we include *Stock*, the monthly level of the S&P 500 index in the employee's first month, because it captures changes in the attractiveness of retirement savings accounts.

5.3 Distributional Analysis

Of particular interest are responses to the increase in mandatory contributions of employees who are likely to save more or less, since the full crowd-out benchmark predicts reductions on the intensive margin for the former and the extensive margin for the latter; and of employees who are more or less likely to face liquidity constraints, and thus experience greater reductions in welfare. To pinpoint these responses, we focus on two features of the data: employees at the low end of the distribution of voluntary contribution rates, who may reveal a preference for either liquidity or commitment, and employees who have relatively low earnings, who may experience greater changes in welfare.

To do this, we undertake quantile analysis, comparing the distributions of voluntary contributions for workers hired before versus after the change in required contributions. We display the raw distributions of voluntary contributions by new hires before versus after the policy change, allowing for a straightforward visual contrast against the full crowd-out benchmark. We then supplement these comparisons with unconditional quantile regressions using the method of Firpo (2007), which enables us to include controls and assess statistical significance. We focus on voluntary contributions at the end of the first academic year of employment (in month nine), when contribution rates have reached a plateau. To maximize power when undertaking this analysis, we split the sample approximately in half, corresponding to an initial salary of $75,000.

6 Results

In this section, we report estimates of changes in voluntary contribution rates, as a percent of salary, for employees hired after the shift in required DC plan contributions, compared to employees hired before. As summarized in Table 2 and Figure 2, our benchmark theoretical model showed that optimizing savers who would save a considerable amount on their own should reduce their voluntary DC contributions by 3.5-5 pp in response to the new 5 pp required employee contribution and 1.5 pp reduction in the required employer contribution; while optimizing savers who would save only a little should stop participating in the voluntary plan.

While we begin with OLS estimates for the full sample, the results of primary interest are the responses along the full distribution of voluntary contributions, for low and high earners. A participation decline in the parts of the distribution where employees were contributing less than 5%, especially among lower earners, reflects a demand for liquidity that can only partially offset the increase in mandatory contributions. In contrast, a lack of responsiveness in voluntary contributions reflects a demand for commitment. After estimating the magnitudes of these responses, we compute the value of liquidity and of commitment for the respective groups.

6.1 Full-Sample Estimates

We present estimation results in Table 4 for voluntary participation and contribution rates for the full sample, based on our event study specification.[27] We report a subset of the coefficient estimates on the interactions of *Post* tenure month, focusing on the 3rd, 6th, and 9th months after hire, and we also report the coefficient estimates on the age and earnings interval dummies (while omitting the additional demographic variables). Figures 4a-4c display the estimates with 95 percent confidence intervals for the full set of first-year-of-hire coefficients.

The change in the overall voluntary participation rate, measured in percentage points, is small in magnitude, not statistically different from zero starting in the 3rd month, and slightly positive, though theory predicts a reduction in participation.[28] Thus, the point estimates hover around zero, and the lower bound of the 95 percent confidence intervals show an average decrease of at most 5pp, and closer to 3 pp. Even this lower bound estimate is far from the potential for crowd-out on the participation margin, since about 50 percent of participants were contributing above zero and below 5% in the pre-treatment period.

We find more evidence of a crowd-out response for the amount contributed within the full sample. The remainder of Table 4 shows the estimated effect of the shift in required contributions on the voluntary contribution rate in the 3rd, 6th, and 9th months after hire, and Figures 4b and 4c show the estimates for each month. Conditional on a positive contribution, new faculty hired after the policy change reduce their average voluntary contribution rate by 2.5-3 pp, a statistically significant reduction.[29] For the full sample including those with zero contributions, the average reduction is generally between 1.5-2 pp.

The evidence of limited crowd-out is supported by an event-study analysis that appears in Figure A.7. While the estimates are somewhat noisy, likely reflecting our small sample

[27]Notably, we find a small and insignificant reduction in the initial one-time choice of the faculty DC plan, relative to the state DB plan, following the simultaneous shift in required DC contributions that we study and the moderate reduction in promised future DB pension wealth. The 2.3 pp reduction, as shown in Appendix Table A.2 is small relative to baseline DC participation of 79%, so we have little ground to be concerned about compositional changes among DC plan participants. As robustness, we also estimate our main results using all faculty, including those who chose the DB plan, and find similar results (Figure A.9 and Table A.6).

[28]In the first two months, participation is substantially higher in the *Post* period, while the contribution rate conditional on participating is quite similar.

[29]We interpret this result descriptively rather than causally since it conditions on making a positive contribution.

size, they show a reduction in average voluntary contributions that is considerably less than 3.5-5pp. The one-year reduction in voluntary contributions of new hires from 2009 to 2010 is most notable, reaching about 2 pp, which perhaps indicates greater salience of the shift in mandatory contributions for new hires in 2010 than for later cohorts of new hires. The long-run reduction from before the Great Recession to the steady state that appears from 2011 on is also about 2 pp. Differences across other intervals before and after 2010 suggest smaller reductions.

The results indicate crowd-out rates far below what is predicted by standard economic theory. The crowd-out should be 3.5-5 pp reduction, yet the lower bounds of the 95 percent confidence intervals reject even a 3.5 pp reduction in contribution. Moreover, the participation response is far too small to explain the lack of crowd-out in the full sample. Overall, we estimate a "crowd-out rate", as in Card and Ransom (2011), of about 30%, compared to their estimate of between 60 and 80% to mandated employee contributions.[30]

6.2 Distributional Analysis by Earnings Level

Now, we use quantile plots to evaluate crowd-out effects along the distribution of voluntary contributions. We display the raw distribution of contribution rates of workers hired before versus after the policy change. We further distinguish between lower and higher earners, using the approximate median initial salary of $75,000. We then supplement these comparisons with quantile regression, which enables us to include controls and assess statistical significance.

Figure 5 displays the distribution of contribution rates for employees hired before 2010 (solid black line) and after (dashed black line), separately for those with lower and higher salaries. As a benchmark, we construct the predicted distribution (dashed red line) under full crowd-out by subtracting 5 percentage points (or the full value of the contribution rate if below 5%) from voluntary contributions of employees hired before 2010. We top-code contributions at 25 percent of salary to show meaningful variation between groups at lower contribution rates.[31]

These quantile plots show a modest participation response among low-contributing low earners, and a modest crowd-out response among high-contributing high earners. Starting with employees with salaries below $75,000, in the left panel of Figure 5, non- participation

[30] Dividing the 1.496pp reduction in voluntary contributions after 9 months from Table 4 by 5pp yields a 30 percent crowd-out rate.

[31] We also examine the probability of making the maximum contribution, which is governed by 403(b) regulations, and find a small reduction of 3 pp (Appendix Figure A.8), which is about 30 percent of the baseline mean.

goes from about 30 percent before the policy to just over 40 percent after. Yet, this is far below the drop in participation expected under the full crowd-out benchmark, in which over 80 percent of lower-earning employees would cease making voluntary contributions. Instead, not only do employees above the 45th quantile continue to participate, but they hardly change their contributions at all, as indicated by a less than 1 pp gap between the solid (pre-2010) and dashed (post-2010) black lines.

Among employees with salaries over $75,000, in the right panel of Figure 5, the patterns are different. A very small decline in participation occurs, even though a drop of over 40 percentage points would be expected under full-crowd out. Instead, the empirical distributions of employees hired before versus after 2010 are nearly identical until around the median contribution rate. Above the median, the distributions diverge, with evidence of full crowd-out: the post-2010 distribution closely tracks the full crowd-out benchmark (dashed red line) of 5 pp.[32]

To formally examine these distributional differences, we estimate unconditional quantile regressions at each percentile of the contribution rate distribution using the method of Firpo (2007). The regression specification mirrors equation (4), including the same covariates, estimated for each quantile q separately for employees below and above the median salary. The coefficient on *Post* in the qth regression indicates the change in y_q, the qth percentile of the unconditional voluntary contribution, between employees hired after 2010 to employees hired before. Figure A.2 presents the coefficient estimates on *Post* from each percentile, with the shaded region representing 95 percent confidence intervals.[33] The vertical line is placed at the quantile that corresponding to a 5-percent contribution rate among faculty hired prior to 2010.

The unconditional quantile regressions confirm the patterns documented in Figure 5. For employees below $75,000 in salary, the changes are not statistically distinguishable from zero across the entire distribution of contributions (Figure A.2a). By contrast, for employees above $75,000 in salary, there are reductions of around 5 percentage points between the 55th and 85th quantiles, as well as at some higher quantiles (Figure A.2b).

In summary, these results provide evidence of incomplete crowd-out for employees. Notably, we find a similar percentage of employees responding with full crowd-out, at around 17%, as the percentage of active savers in the analysis of mandatory contributions in Denmark

[32]Notably, crowd-out for this group reaches 5 pp, the upper bound of the theoretically predicted range, rather than the lower bound of 3.5 pp

[33]For visual clarity, we present results from the 45th to 93rd percentiles. In the bottom half of the distribution, there is either mechanically no change in contributions (when participation is zero) or no meaningful difference. At the very top of the distribution, the estimates are noisy and the wide confidence intervals prevent important variation for the rest of the distribution from being clearly shown on the same scale.

in Chetty et al. (2014). Moreover, we demonstrate important heterogeneity in responses for low and high earners. The policy change produces moderate reductions in the probability of participating among lower-salaried employees, though theory predicts participation responses that are four times larger than observed. For other employees with low salary, there is little change in the level of contributions. Among those with higher salaries, there is evidence of crowd-out at the higher end of the contribution range but no change in participation rates.

6.3 Robustness

We implement several robustness checks and find similar results: we continue to strongly reject full crowd-out. First, we estimate specification (4) without any of the control variables for demographics, job characteristics, and financial market returns. Appendix Figure A.3 shows that the response is slightly weaker, with most estimates around a 1 pp reduction in contribution rates, and lower bounds of the 95% confidence interval above a 3 pp reduction. We then return to the original specification but omit employees hired during the financial crash, from September 2008 through end-2009, because financial choices may have changed during this window. The main estimates of crowd-out become noisier but also larger in absolute value, though we might have expected the opposite. If anything, contributions might have been suppressed for employees hired during the financial crash and boosted during the later recovery, which would lead to an underestimate of crowd-out. In any case, we continue to reject a 5 pp reduction based on the lower bound of the 95% confidence intervals (Appendix Figure A.4). Next, we extend the analysis to the first 3 years since the employee was hired. While we earlier noted some compositional differences for faculty who remain employed after 2 years, the events studies in Appendix Figure A.5 show similar patterns for the first 27 months as the first 9 months.

Next, we perform a placebo test with our main sample using the choice of plan vendor that employees enroll with as the dependent variable. The choice between two vendors is a decision made concurrently with first choosing the contribution rate, but should arguably not be affected by the mandatory increase in employee contributions. If other unobservable characteristics of employees who join the university changed over this period, or if changes in financial markets or the macroeconomy had occurred at the same time, then this choice might change as well. Appendix Figure A.6 shows that estimates of the policy change impact on the choice of vendor are very close to zero and are not statistically significant.

Finally, we leverage the fact that staff were largely unaffected by the policy change to conduct a difference-in-difference analysis. Using staff as a control group helps to address concerns that the results may be due to changes across cohorts in preferences, in composition, or in information about retirement saving, whether driven by changes at the employer or in the economy.[34] We find essentially no change in the staff voluntary contribution rate, either in 2010 or gradually over time. Appendix Figure A.7 plots average contribution rates separately for faculty and staff, and split by whether the employee was hired before or after the policy change. The figure plots rates that have been regression-adjusted using the S&P 500 level in that calendar month and flexible controls for age, income, gender, and full-time status. Contribution rates for staff are quite flat over time, and do not exhibit a break in 2010.[35] By contrast, contribution rates for faculty fall among employees hired after the policy change. Appendix Table A.6 presents the difference-in-difference version of this figure and estimates a 1.0 percentage point decline in contribution rates. The lower bound of the 95% confidence interval is a -1.6pp reduction, which is approximately the event-study magnitude from Table 4, that compares successive cohorts of faculty. Adding staff also increases precision of the crowd-out estimate. In summary, using staff as a control group provides further empirical support to our finding that crowd-out is limited, on average, and is well below the neoclassical benchmark.

7 Welfare Effects of Mandatory Contributions

This section shows how we weigh the costs and benefits of mandatory contributions for different groups of employees. Under the benchmark model of optimizing savers presented in Section 3, the policy does not affect utility for employees already saving above 5%, who can offset the mandatory increase through reduced voluntary contributions, while it harms employees employees with lower saving rates who must as a consequence reduce current consumption. Yet, the quantile plots in Section 6 show a lack of crowd-out for a majority of employees. This example of "passive saving" rather than "active saving" (terms used in Chetty et al. (2014)) calls into question whether individuals were saving optimally; if not, then they may be better off under mandatory contributions. The optimal contribution policy in models by Beshears et al. (2020) and Moser and Olea de Souza e Silva (2019) trades off the future consumption increases for passive savers who otherwise consume too much today against the current consumption declines for active savers who need liquidity. We use stylized calculations of consumption-equivalents to evaluate this trade-off.

We begin by classifying employees into four groups based on their estimated responses to the policy. Each is associated with different welfare effects:

[34] The changes in the state DB plan that we noted earlier resulted in modest reductions in promised DB pension wealth at retirement age, which might lead to more voluntary contributions through a crowd-in effect or less contributions through a reduced-wealth effect.

[35] Appendix Table A.5 shows summary statistics for staff split by their date of hire. Gender and age are similar, while more employees report their marital status after 2010.

1. *Active high savers*: employees saving more than the mandated increase who respond with complete crowd-out; their welfare is unaffected by the policy.

2. *Active low savers*: employees saving less than the mandated increase who stop making voluntary contributions; they value liquidity and are harmed by the policy. Their welfare loss is limited, nevertheless, as their prior willingness to make a voluntary contribution makes part of the crowd-out inframarginal, while the rest reflects a new constraint on their liquidity that is not very binding.

3. *Passive savers*: employees who display little or no crowd-out and continue making voluntary contributions; they value commitment and are helped by the policy. Here, we assume that the new required contribution brings them to their optimal consumption level, which results, conservatively, in a relatively small welfare gain.

4. *Non-savers*: employees who did not make voluntary contributions in the first place; since no actions reveal their type, we consider alternately that they value liquidity or value commitment. In either case, we make the conservative assumption that either they were consuming optimally beforehand and did not face a prior liquidity constraint (or else their welfare loss will be bigger), or they are now saving optimally and face no further commitment failure (or else their welfare gain will be bigger).

Active high savers: In the theory we laid out earlier, active high savers experience no change in well-being, so they drop out of the welfare calculations. In Figure A.2b, 35 percent of employees with salaries over $75,000 exhibit full crowd-out.[36] We fail to find evidence of full (or any) crowd-out among employees with salaries below $75,000 except on the participation margin, and therefore estimate 17.5 percent of employees in our sample to be active savers.

Active low savers: We measure the cost of mandatory contributions to active low savers using a consumption-equivalent approach. The consumption-equivalent measures the dollar amount a person would be willing to forgo each period to avoid making mandatory contributions while working. It reflects the cost of distorting the optimal consumption profile by shifting more lifetime consumption into the future and away from the present. To estimate this cost, we calculate the average consumption equivalent for active low savers using our two-period model from Section 3. We parameterize utility as constant relative risk aversion,

[36]The 35 percent corresponds to quantiles 60 to 99, excluding quantiles 84 to 88, which have point estimates between -1.0 and -3.4.

with $u(c) = \frac{c^{1-\eta}}{1-\eta}$ and η denoting the coefficient of relative risk aversion.[37] For each active low saver, the consumption equivalent is the amount of money x that solves the following equation:

$$u(c_1^*(1-x)) + \delta u(c_2^*(1-x)) = u(c_1^* - M_f) + \delta u(c_2^* + (1+r)M_f) \quad (5)$$

where M_f denotes the mandatory contribution in dollar terms that cannot be offset by lowering voluntary contributions (based on the pre-treatment range of voluntary contribution amounts that get crowded out) and optimal consumption in each period is denoted by c_1^* and c_2^*. Given non-zero voluntary contributions that these employees made before the policy change, which we treat as revealing optimal saving, we calculate optimal consumption in the first period as the employee's salary less the full voluntary contributions among those employees hired before the policy change: $c_1^* = y - M_v$, with M_v denoting voluntary contributions. We then calculate c_2^* from the Euler equation: $u^{'}(c_1^*) = \delta u^{'}(c_2^*)$. We calculate x assuming CRRA coefficients equal to 3 or 4, and assuming $\delta = 0.97$ and $r = 0.03$.

As shown in Figure 5 and in Table 5, we estimate an 11 percentage point decrease in the probability of voluntary participation among those below \$75,000, and a 4 percentage point decrease among those over \$75,000, yielding 7.5% active low savers overall. Table 6 presents their average-consumption equivalent, with the upper-bound column excluding non-savers from this calculation by assuming that they value commitment rather than liquidity. The average consumption-equivalent per year with a CRRA coefficient of 3 is \$94 per employee. With a CRRA coefficient of 4, which implies a greater value of consumption-smoothing, it rises by almost half, to \$135. The relatively low values reflect the observation that the active low savers were making voluntary contributions beforehand and hence the new constraint on their liquidity neither is large nor pulls them far from their optimal consumption level. One way to benchmark these magnitudes is in terms of annual consumption. At the higher level of risk aversion, the consumption-equivalents are between 0.3% and 0.5% of annual consumption, on average.

Passive savers: Based on the bounds on active savers provided above, we estimate that passive savers, whose voluntary contributions remain largely unchanged, comprise 50% of the sample; this ignores for now the possibility that non-savers are also passive in their

[37] A richer life-cycle model would include additional periods; stochastic labor income, financial market, and mortality processes;and additional motives for saving. This would involve many additional assumptions, but is likely to deliver substantively different results only in the cases below of passive savers and non-savers, for whom we cannot say how far individuals are from their optimal saving levels, and therefore how much farther (for those who value liquidity) or closer (for those who value commitment) the mandatory contribution takes them.

savings decisions. A large body of evidence points to employees likely under-saving relative to their optimum, rather than over-saving. A number of mechanisms may explain such under-saving, including time-inconsistency, salience, inattention, incorrect beliefs about longevity or health care costs, inertia, and financial illiteracy (Beshears et al. 2018).[38] Characterizing the exact mechanisms in our setting is not of primary interest, however, because it does not affect the welfare gains to passive savers: mandatory contributions reallocate consumption towards later periods, where it is needed, regardless of the underlying mechanism behind their initial choice.[39]

We estimate the welfare gains to passive savers using a symmetric consumption-equivalent approach to the case of active low savers. The perspective now shifts to viewing the saving decisions under mandatory contributions as moving passive savers towards their optimal level. Unlike the case of active low savers, however, where the prior voluntary contribution amount pins down the change in liquidity that they experienced, we cannot determine how close passive savers get to their optimum (or, in fact, whether they overshoot it). Therefore, we make the relatively conservative assumption that the mandatory contribution moves them exactly to their optimum; if they remained below the optimum, the welfare gain would be higher than what we calculate.[40] Thus, optimal consumption in period 2 is given by salary less the sum of mandatory contributions and any voluntary contributions among employees hired after 2010: $c_2^* = y - M_f - M_v$. Denoting β as the amount of mandatory contributions that are crowded-out, the consumption equivalent for passive savers is thus the amount z that equates:

$$u(c_1^*(1-z)) + \delta u(c_2^*(1-z)) = u(c_1^* + M_f + \beta) + \delta u(c_2^* - (1+r)(M_f + \beta)) \tag{6}$$

We calculate the average z among passive savers using the post-2010 distribution of contribution rates, salaries, and the quantile-specific estimates of β from Figures A.2a and

[38] If passive savers were instead saving too much, on average, then forcing them to save even more only further distorts their consumption profile and the policy is unambiguously welfare-reducing. We therefore focus our attention on the case that passive savers are under-saving since that represents a potential trade-off in policy design.

[39] As discussed in Handel and Schwartzstein (2018), pinpointing the reasons for sub-optimal behavior are not necessary for welfare evaluation of policies that force all consumers to make the same decision.

[40] If, on the other hand, mandatory contributions lead workers to overshoot the optimum, then the gain would be lower. To assess whether passive savers are likely below or above the optimum, we calculate the increased annuity amount at age 65 implied by the increased saving rates due to mandatory contributions. On average, the annuity amounts to about 6.7 percent of final projected salary. The increase is larger for younger workers, who have more years of saving relative to retirement. This represents a meaningful increase, but still appears to leave many workers who are not young below target levels of saving at retirement according to calculations by Munnell et al. (2019).

A.2b.

Table 6 presents the average consumption-equivalent among employees classified as passive savers, with the lower-bound column not including non-savers, under the assumption that they value liquidity rather than commitment. The average welfare gain for passive savers is $202 when the CRRA coefficient is 3 and $258 when it is 4 under the upper bound assumption. This exceeds the average welfare loss for active low savers because passive savers have higher salaries, so the change in saving for them is greater, and because the full amount of extra saving is assumed to improve welfare. As a fraction of consumption, the consumption equivalents are slightly greater than 0.5% for passive savers.

Non-savers: Our data do not allow us to discern whether employees who did not make voluntary contributions before the policy actively chose not to (perhaps because of liquidity constraints), or whether they experience greater commitment failures. We therefore assume, as an upper bound on passive savers, that the required contribution also brings them to their optimal savings level; and as a lower bound, that the required contribution moves them away from their optimal savings level by constraining their liquidity.

Initial non-savers constitute 33% of lower-salaried employees and 18% of higher-salaried employees, or about 25% of the full sample. Including them in the group of passive savers raises their share of the sample to 73% and yields nearly the same mean consumption-equivalent of their welfare gain ($203 vs $202) when the CRRA coefficient is 3. Including them instead in the group of active low savers group instead raises their share to 33% of the sample, and substantially raises the mean consumption-equivalent of the welfare loss from $94 to $170, as many of the non-savers are at the low end of the salary distribution.

Overall costs and benefits: With passive savers making up a majority of the sample, even under the conservative assumption that all non-savers value liquidity rather than commitment, we find that the aggregate welfare gains exceed the costs to low active savers in our setting. Under this conservative assumption, gains exceed losses by a ratio of $1.6 - 1.9$, while the ratio climbs to $19 - 21$ with an upper bound assumption that all non-savers value commitment. Appendix Table A.7 shows slightly higher ratios of benefits to costs when we just consider employees with salaries below $75,000. The qualitative findings are thus not driven solely by higher-salaried employees who have greater willingness-to-pay for a smooth consumption profile. The incidence of the policy, however, is important to highlight: the costs of this policy, ignoring the benefit side, are concentrated among lower-salaried employees.

These calculations omit several factors that are worth noting. First, we are unable to observe whether employees reduce other forms of saving outside employer accounts. Should that occur, the crowd-out rate would be higher, and our methods would over-estimate the fraction of passive savers. However, it seems likely that employees would first adjust within their DC plans. Second, we have assumed all passive savers were under-saving, rather than some over-saving, before mandatory contributions were enacted. Third, borrowing is allowed against assets in voluntary accounts, which can alleviate liquidity constraints; in our setting, though, the incidence of borrowing by participants against DC plan balances is quite infrequent among those who do participate, with less than 2% taking out any loans.

8 Discussion

Even though the goal of requiring employees to contribute to a DC pension plan is presumably to induce more saving, standard economic theory predicts crowd-out, as higher required contributions should lead employees to reduce their voluntary contributions. Higher total saving would emerge among employees who would have saved little on their own, but at the cost of reducing their welfare by absorbing some of their needed liquidity. Our findings in this paper reject the predictions of standard theory. Comparing new employees hired after a large public university increased required employee contributions by 5 pp, while reducing the employer contribution by 1.5 pp, to new employees hired before, we detect reductions of voluntary contribution rates of around 1.5 pp, on average, statistically rejecting the 3.5-5pp reductions predicted by theory.

The passive saving behavior that we observe supports evidence from a variety of settings that individuals do not fully optimize over their retirement savings choices. Further evidence that default enrollment fails to raise savings rates over the long run has turned attention to stronger paternalism mechanisms like required contributions. The degree to which such mechanisms can make people better off depends on the trade-off between liquidity, by those who are optimally saving very little, and commitment, by those who would benefit from saving more. We use our evidence of incomplete crowd-out to assess this trade-off.

In order to characterize the responses that we observe, we consider where crowd-out originates along the distribution of contribution rates for high- versus low-salaried employees in our sample. We observe distinct patterns, with full crowd-out among high earners at the high end of the distribution and among low earners at the low end of the distribution of voluntary contribution rates.

Given these findings, we map the shares of employees with differing crowd-out responses into groups based on their apparent value for liquidity versus commitment, and we apply a stylized model to calculate the implied welfare losses versus gains. We find the following: (1) 17.5% are unaffected by the policy, based on where full crowd-out emerges at the high end of the distribution of contribution rates. Their total DC plan saving remains approximately unchanged. (2) At least 7.5% (and as much as 33%) appear to value liquidity, based on the reduced participation rate in the voluntary plan (or also counting initial non-savers before the policy change). This group cannot fully crowd-out required contributions, and they experience an average welfare loss of $94 (or $170 under the more inclusive assumption). (3) Lastly, at least 50% (and as much as 75%) appear to value commitment, based on the lack of crowd-out for remaining contributors (or also counting the initial non-savers, who may be passive savers). This group increases their total DC plan saving and experiences an average welfare gain of about $200.

Because the group of passive savers is so large, the value of their gains exceeds the losses to active savers who value liquidity. Even under treating all initial non-savers as valuing liquidity, the gains to those who value commitment are 1.6 times higher than the losses to those who value liquidity. In sum, our results show the effectiveness of using required DC plan contributions to raise savings rates and, by showing who reacts and who does not, helps weigh the costs and benefits of such a policy.

It is important to acknowledge our results pertain to a single setting, and the employees we study have relatively high incomes and education levels. The costs of mandated retirement saving may be larger in other settings with a greater share of active low savers who are liquidity constrained. Yet, if passive savers make up the majority of employees in other settings as they do in ours, then the gains from commitment may nonetheless outweigh the costs. Employers and policymakers should weigh these trade-offs when considering whether defaults should be replaced or combined with more paternalistic policies to increase retirement saving.

References

Amador, Manuel, Ivan Werning, and George-Marios Angeletos, "Commitment vs. Flexibility," *Econometrica*, 2006, *74* (2), 365–396.

Beshears, John, James Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte Madrian, "Optimal Illiquidity," 2020.

_ , _ , David Laibson, and Brigitte Madrian, "Behavioral Household Finance," in Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *Handbook of Behavioral Economics*, Stanford University Press, 2018, pp. 177–276.

_ , _ , _ , _ , and William Skimmyhorn, "Borrowing to Save? The Impact of Automatic Enrollment on Debt," *Journal of Finance*, forthcoming.

Bubb, Ryan and Patrick Warren, "An Equilibrium Theory of Retirement Plan Design," *American Economic Journal: Economic Policy*, 2020, *12* (2), 22–45.

Buchanan, James, "The Samaritan's Dilemma," in Edmund Phelps, ed., *Altruism, Morality, and Economic Theory*, Russell Sage Foundation, 1975, pp. 71–85.

Card, David and Michael Ransom, "Pension Plan Characteristics and Framing Effects in Employee Savings Behavior," *Review of Economics and Statistics*, 2011, *93* (1), 228–243.

Carroll, Gabriel, James Choi, David Laibson, Brigitte Madrian, and Andrew Metrick, "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics*, 2009, *124* (4), 1639–1674.

Chetty, Raj, John N Friedman, Soren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen, "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark," *The Quarterly Journal of Economics*, 2014, *129* (3), 1141–1219.

Choukhmane, Taha, "Default Options and Retirement Saving Dynamics," 2019.

Clark, Jeffrey and Jean Young, "Automatic Enrollment: The Power of the Default," 2018.

Feldstein, Martin, "Rethinking Social Insurance," *American Economic Review*, 2005, *95* (1), 1–24.

Firpo, Sergio, "Efficient Semiparametric Estimation of Quantile Treatment Effects," *Econometrica*, 2007, *75* (1), 259–276.

Friedberg, Leora and Adam Leive, "Overpaying and Undersaving: Correlated Mistakes in Retirement Saving and Health Insurance Choices," 2021.

_ and Michael Owyang, "The Rise of 401(k) and Other Defined Contribution Plans.," *Federal Reserve Bank of St. Louis Review*, 2002, *84* (1), 23–34.

Handel, Benjamin and Josh Schwartzstein, "Frictions or Mental Gaps: What's Behind the Information We (Don't) Use and When Do We Care?," *Journal of Economic Perspectives*, 2018, *32.*

Kaplan, Greg, Giovanni Violante, and Justin Wiedner, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, pp. 77–138.

Laibson, David, "Nudges are Not Enough: The Case for Price-Based Paternalism," 2020.

Madrian, Brigitte C and Dennis F Shea, "The power of suggestion: Inertia in 401 (k) participation and savings behavior," *The Quarterly journal of economics*, 2001, *116* (4), 1149–1187.

Moser, Christian and Pedro Olea de Souza e Silva, "Optimal Paternalistic Savings Policies," 2019.

Munnell, Alicia H, Wenliang Hou, and Geoffrey T Sanzenbacher, "How Would More Saving Affect the National Retirement Risk Index," 2019.

Price, T. Rowe, "Reference Point," 2020.

Saad, Lydua, "401(k) Regaining Importance as Future Income Source," May 2017.

Thaler, Richard and Cass Sunstein, *Nudge: Improving Decisions about Health, Wealth, and Happiness.*, Yale University Press, 2008.

_ **and Shlomo Benartzi**, "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, 2004, *112* (S1), S164–S187.

Table 1: Summary of Policy Change to Contribution Rates, percent of salary

	Hired Before July 1, 2010	Hired After July 1, 2010
Mandated Employee Contribution Rate	0%	5%
Employer Contribution Rate	10.4%	8.9%
Total Contribution Rate	10.4%	13.9%

Note: This table summarizes the changes in the university DC plan based on the employee's date of hire. Employees faced a one-time irrevocable choice between a state DB plan and a university DC plan. The changes were the result of state law enacted to alleviate underfunding in the state DB plan.

Table 2: Summary of Model Predictions

	Average voluntary contribution rate	Average voluntary participation rate
Baseline case (*with crowd-out and compensation reduction effects)*	Decrease between 3.5 and 5pp	No change
Low saving rate case (*with additional participation effect*)	Smaller decrease than above	Decrease

Note: This table summarizes the model's predictions for participation and voluntary contributions in response to the required increase in contributions in the university DC plan. Under the baseline case, employees offset the required 5 percent increase in mandatory contributions by reducing their contributions between 3.5pp and 5pp. The 5 percent reduction represents the crowd-out effect in which lower employee contributions completely offset the 5% increase in required contributions. The reduction could be as low as 3.5pp due to the compensation effect, though this effect is likely to be small if that the 1.5pp reduction is scaled by the average contribution rate of 6-8%. For employees who contribute small enough amounts that they cannot fully reduce their contributions, they are predicted to stop participating by making zero contributions, reflecting a participation effect.

Table 3: Descriptive Statistics

	Pre: Hired 2005-2009		Post: Hired 2010-2014		Diff.	p-val
	Mean	SD	Mean	SD		
Choice between the DB plan or mandatory DC plan:						
DB plan participation	0.210	0.407	0.194	0.396	-0.015	0.332
Mandatory DC plan participation	0.790	0.407	0.806	0.396	0.015	0.332
If in the mandatory DC plan:						
Voluntary DC participation rate	0.761	0.427	0.761	0.427	0.001	0.977
Voluntary DC contribution rate, %	5.231	8.987	4.437	8.600	-0.793	0.040
Voluntary DC contribution rate, % (if positive)	7.273	11.754	6.070	10.665	-1.203	0.035
Demographic and job characteristics:						
Female	0.429	0.495	0.432	0.496	0.003	0.877
Married	0.439	0.496	0.438	0.496	0.000	0.995
Single	0.243	0.429	0.258	0.438	0.015	0.375
Marital status unknown	0.319	0.466	0.304	0.460	-0.015	0.412
Full-time	0.911	0.285	0.908	0.289	-0.003	0.806
Annual salary	84,373	43,283	87,473	44,656	3,100	0.106
Age	39.73	9.07	38.92	9.14	-0.81	0.023
S&P 500 Index	1261	225	1737	285	476	0.000
p-value from omnibus balance test						0.513
N	1,808		1,024			

Administrative data on faculty at a large public university, statistics from their ninth month of hire. The final column lists the p-value from the t-test the means are equal between employees hired in 2005-2009 vs. employees hired in 2010-2014. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each. The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary. Age and annual salary are reported in bins to preserve confidentiality of individuals in the data set. Age is grouped into bands, generally of five years, and we impute individual age as the mid-point of the age bands. We impute salary as the midpoint of salary bands, winsorized at the 5th and 95th percentiles to remove the influence of outliers on the mean. Both salary and the S&P index are inflation-adjusted to 2018. The omnibus balance test includes demographics and job characteristics and excludes the the S&P index.

Table 4: OLS Regression Estimates, Voluntary Participation and Contributions

	Voluntary Participation		Voluntary Contribution Rate (if pos.)		Voluntary Contribution Rate	
	Estimate	SE	Estimate	SE	Estimate	SE
Post-2010 x Tenure month 3	0.00639	(0.0249)	-2.711	(0.731)	-1.698	(0.532)
Post-2010 x Tenure month 6	0.00897	(0.0257)	-2.674	(0.739)	-1.652	(0.553)
Post-2010 x Tenure month 9	0.0131	(0.0262)	-2.455	(0.763)	-1.496	(0.582)
Age, relative to $[25, 30)$						
$[30, 35)$	-0.0324	(0.0326)	0.458	(0.745)	0.000500	(0.523)
$[35, 40)$	-0.0430	(0.0342)	1.107	(0.812)	0.424	(0.578)
$[40, 45)$	-0.0816	(0.0365)	1.502	(0.917)	0.499	(0.654)
$[45, 50)$	-0.132	(0.0419)	2.823	(1.175)	1.027	(0.822)
$[50, 55)$	-0.115	(0.0460)	6.315	(1.731)	3.550	(1.248)
$[55, 60)$	-0.0866	(0.0491)	2.342	(1.489)	0.787	(1.077)
$[60+)$	-0.226	(0.0578)	1.404	(1.796)	-0.744	(1.164)
Salary, relative to $[20k - 25k)$						
$[25k - 29k)$	0.171	(0.0610)	0.823	(1.659)	1.651	(0.856)
$[30k - 35k)$	0.0936	(0.0635)	0.508	(1.518)	1.167	(0.722)
$[35k - 39k)$	0.173	(0.0629)	-1.256	(1.392)	0.368	(0.686)
$[40k - 45k)$	0.193	(0.0642)	1.794	(1.619)	2.208	(0.866)
$[45k - 49k)$	0.278	(0.0585)	0.720	(1.524)	2.158	(0.914)
$[50k - 55k)$	0.284	(0.0582)	0.635	(1.506)	2.305	(0.906)
$[55k - 59k)$	0.323	(0.0624)	1.777	(1.706)	3.231	(1.069)
$[60k - 65k)$	0.380	(0.0577)	1.497	(1.656)	3.315	(1.062)
$[65k - 69k)$	0.396	(0.0588)	2.063	(1.564)	3.790	(1.006)
$[70k - 75k)$	0.390	(0.0602)	-0.390	(1.359)	2.049	(0.782)
$[75k - 79k)$	0.384	(0.0664)	0.968	(1.907)	3.050	(1.262)
$[80k - 85k)$	0.349	(0.0695)	2.885	(2.013)	4.088	(1.407)
$[85k - 89k)$	0.471	(0.0602)	1.217	(1.627)	3.681	(1.113)
$[90k - 95k)$	0.393	(0.0628)	2.638	(1.733)	4.190	(1.102)
$[95k - 99k)$	0.419	(0.0624)	2.097	(1.764)	3.653	(1.169)
$[100k - 120k)$	0.472	(0.0495)	5.207	(1.362)	6.849	(0.773)
$[120k - 140k)$	0.438	(0.0602)	4.233	(1.813)	5.880	(1.293)
$[140k+)$	0.441	(0.0550)	-0.363	(1.370)	3.510	(0.768)
N	17961		12725		17961	

Note: Table shows regression results of equation 4 using administrative data on faculty at a large public university. Coefficient estimates (with standard errors clustered by employee). Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each; The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary. All regressions include indicators for post-2010, gender, fulltime employee, marital status, the level of the SP 500 on the 1st of each month, and a constant.

Table 5: OLS Regression Estimates, Voluntary Participation and Contributions

	Voluntary Participation Rate	Voluntary Contribution Rate (if pos.)	Voluntary Contribution Rate
Mean of dep. var. pre-2010	0.769	7.059	5.432
Panel A. Heterogeneity by income			
(Income < median, 75k)×*Post*	-0.116	2.277	1.606
	(0.0438)	(1.375)	(0.962)
N	1935	1485	1935
R^2	0.0909	0.0678	0.0450
Panel B. Heterogeneity by age			
(Age < 35)×*Post*	-0.168	0.588	-0.392
	(0.0414)	(1.193)	(0.914)
N	1935	1485	1935
R^2	0.126	0.0678	0.0615
Panel C. Heterogeneity by gender			
Female×*Post*	-0.0399	1.456	1.098
	(0.0406)	(1.184)	(0.940)
N	1935	1485	1935
R^2	0.124	0.0836	0.0714

Notes: Table shows regression results of equation 4 using administrative data on faculty at a large public university. Coefficient estimates (with standard errors clustered by employee). Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each; The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary. All regressions include indicators for post-2010, gender, fulltime employee, marital status, the level of the S&P 500 on the 1st of each month, and a constant. Regressions in Panel A also include an indicator for income below $75k and age categories. Regressions in Panel B also include indicators for age below 40 and for salary bin. Regressions in Panel C also include indicators for age and salary bins. We choose age 40 and salary of $75k to split the sample as these levels correspond approximately to the median.

Table 6: Welfare Analysis: Consumption Equivalents for Active Low Savers and Passive Savers

	Lower Bound on Passive Savers	Upper Bound on Passive Savers
Panel A. Percentage of each type		
Passive savers	49.5	75
Active low savers	33	7.5
Active high savers	17.5	17.5
Panel B. Costs and Benefits (\$/year), CRRA = 3		
Cost: Avg. consumption equivalent, active low savers	170	94
Benefit: Avg. consumption equivalent, passive savers	203	202
Ratio of benefits to costs, weighted by shares in (A)	1.82	21.52
Panel C. Costs and Benefits (\$/year), CRRA = 4		
Cost: Avg. consumption equivalent, active low savers	239	135
Benefit: Avg. consumption equivalent, passive savers	259	258
Ratio of benefits to costs, weighted by shares in (A)	1.62	19.10

Notes: Table summarizes results of consumption equivalent calculations for different types of savers. Panel A presents the percentage of savers classified according to three types based on the quantile plots from Figure 5. The first column considers the lower bound on the fraction of the sample who are passive savers, which is equivalently the upper bound on who are active low savers. The second column considers the upper bound on passive savers, equivalently the lower bound on active low savers. Panels B and C show the average certainty equivalents for active low savers (costs) and passive savers (benefits) for CRRA coefficients of 3 and 4, respectively. The ratio of benefits to costs is calculated as the average benefits weighted by the percentages in Panel A of each type.

Figure 1: Flow Chart of Retirement Saving Choices

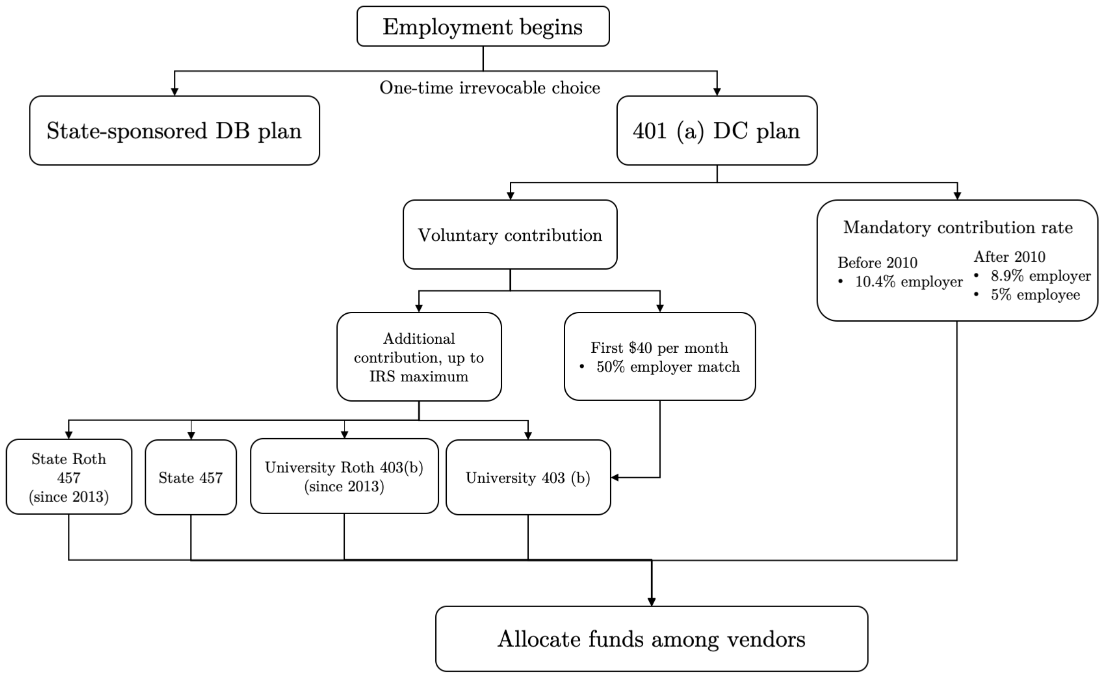

Notes: Figure displays the sequence of choices and available options for retirement plans at the university. The first choice is between the state-sponsored DB plan and the 401(a) DC plan. The main policy of interest is the change in the mandatory contribution rate in the 401(a) DC plan that occurred in 2010. Faculty can choose voluntary contributions with a very limited match rate, of 50% for up to $40 per month, which go into the university-run 403(b) plan. Faculty can choose additional voluntary contributions, which can be directed to the same 403(b) plan and to a state-run 457 plan, and, near the end of our sample period, to a university-run Roth 403(b) plan and a state-run Roth 457 plan. In the last step, employees have choices between two vendors and then numerous plans within each vendor.

Figure 2: Graphical Representation of Theoretical Model

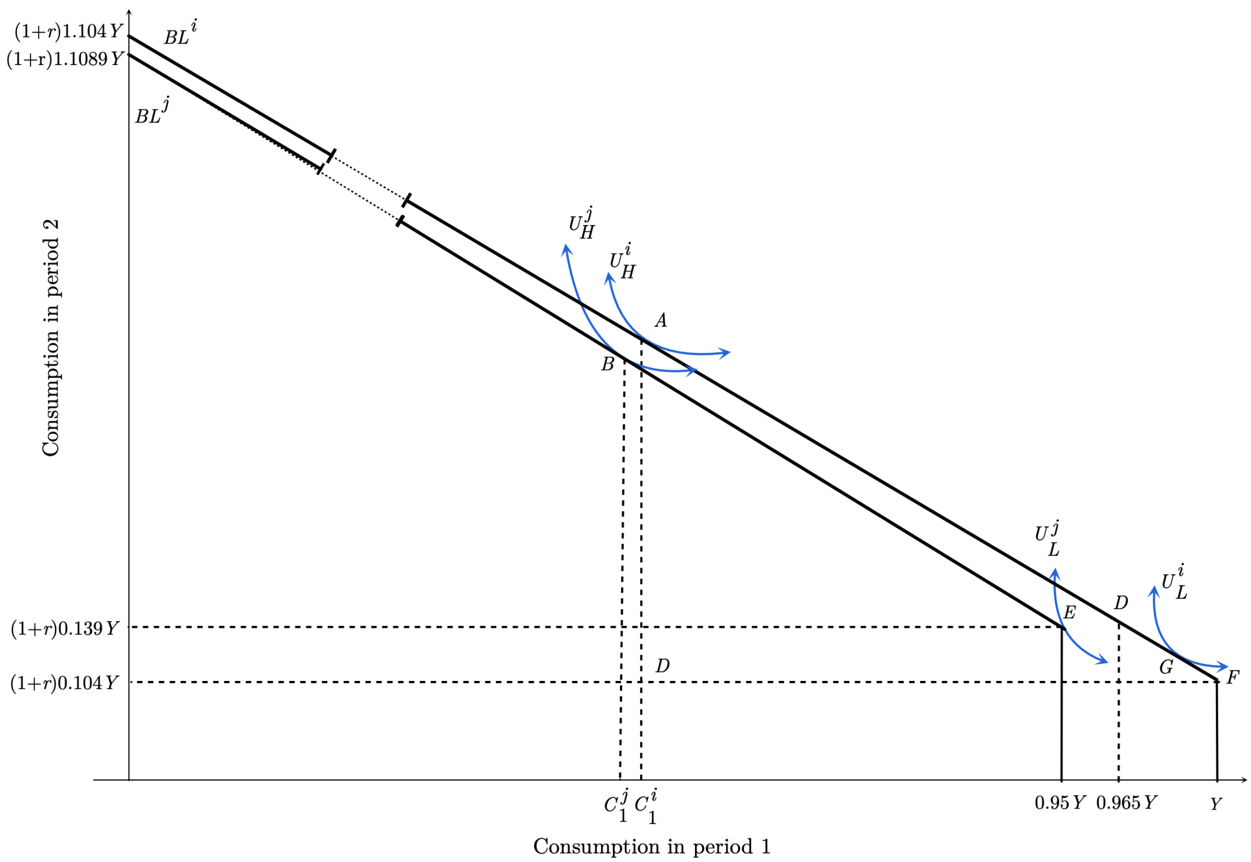

The dotted lines in the budget constraints represent extensions of each line so that the scale is easier to read.

Notes: This graph sketches the predictions of the standard theory of retirement savings crowd-out. The budget line before the policy change (without mandatory contributions) is BL^i and the budget line after the change is shifted inwards to BL^j. Point A represents the optimal consumption bundle for employees hired before the policy change. Point B represents the optimal consumption bundle for employees with the same preferences hired after the policy change, who are now subject to a 5 percent mandatory employee contribution and receive 8.9% in employer contributions, down from 10.4%. Both mandatory and voluntary saving are assumed to have the same rate of return r. The change in voluntary contributions is calculated as $(Y - c_1^i) - (0.95Y - c_1^j) = 0.05Y - (c_1^i - c_1^j)$. This is composed of the sum of a crowd-out effect (equal to -0.05Y) and a compensation effect (equal to $c_1^i - c_1^j$). Point G represents the optimal choice before the policy for an employee with a low saving rate. After the policy change, the employee stops participating, reducing voluntary contributions to zero, and locates at the notch in BL^j at E.

Figure 3: Voluntary Participation and Contribution Rates by Tenure Month, Newly Hired Faculty

(a) Voluntary participation rate

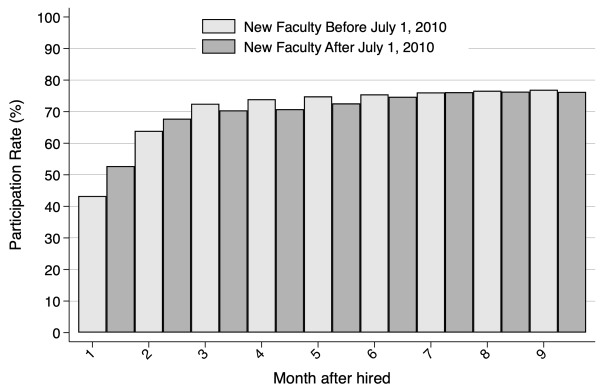

(b) Voluntary contribution rate, if positive

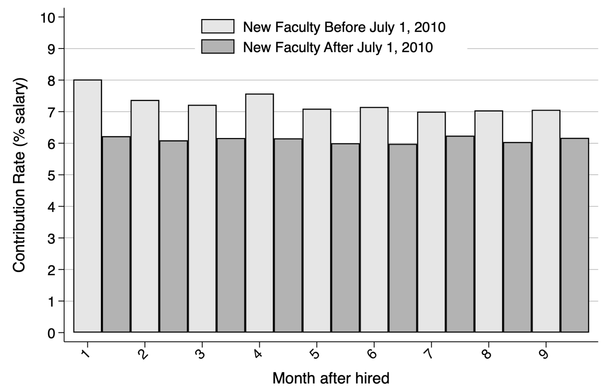

(c) Voluntary contribution rate

Notes: Figure plots mean participation rates and contribution rates by month since the month faculty are hired. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty do not receive a salary. Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each; The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary.

Figure 4: Event study regression estimates

(a) Voluntary participation rate

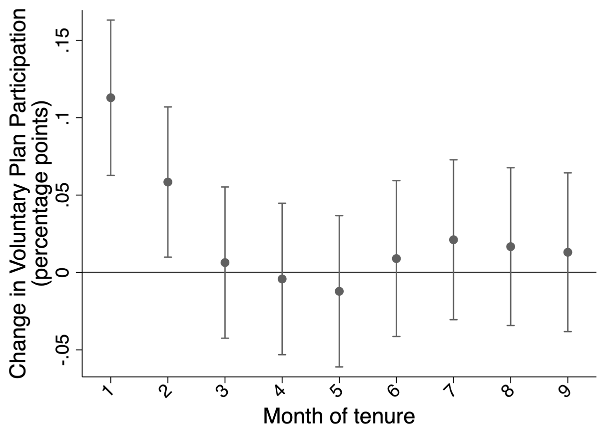

(b) Voluntary contribution rate, if positive

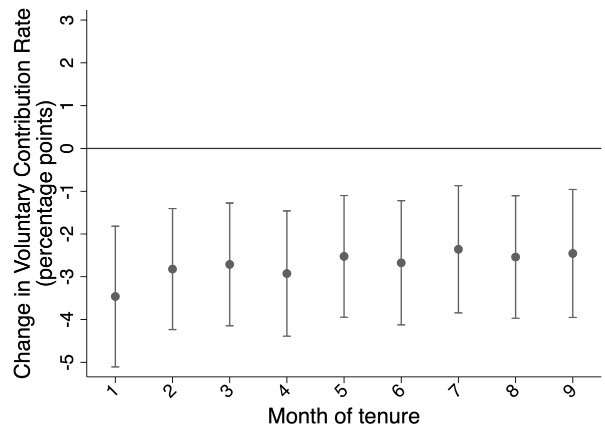

(c) Voluntary contribution rate

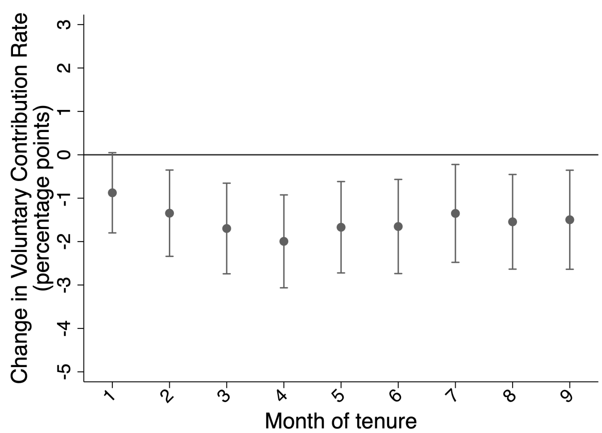

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation 4. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. Additional control variables include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty do not receive a salary. Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each. The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary.

Figure 5: Quantile plot of voluntary contribution rates, by salary level and date of hire

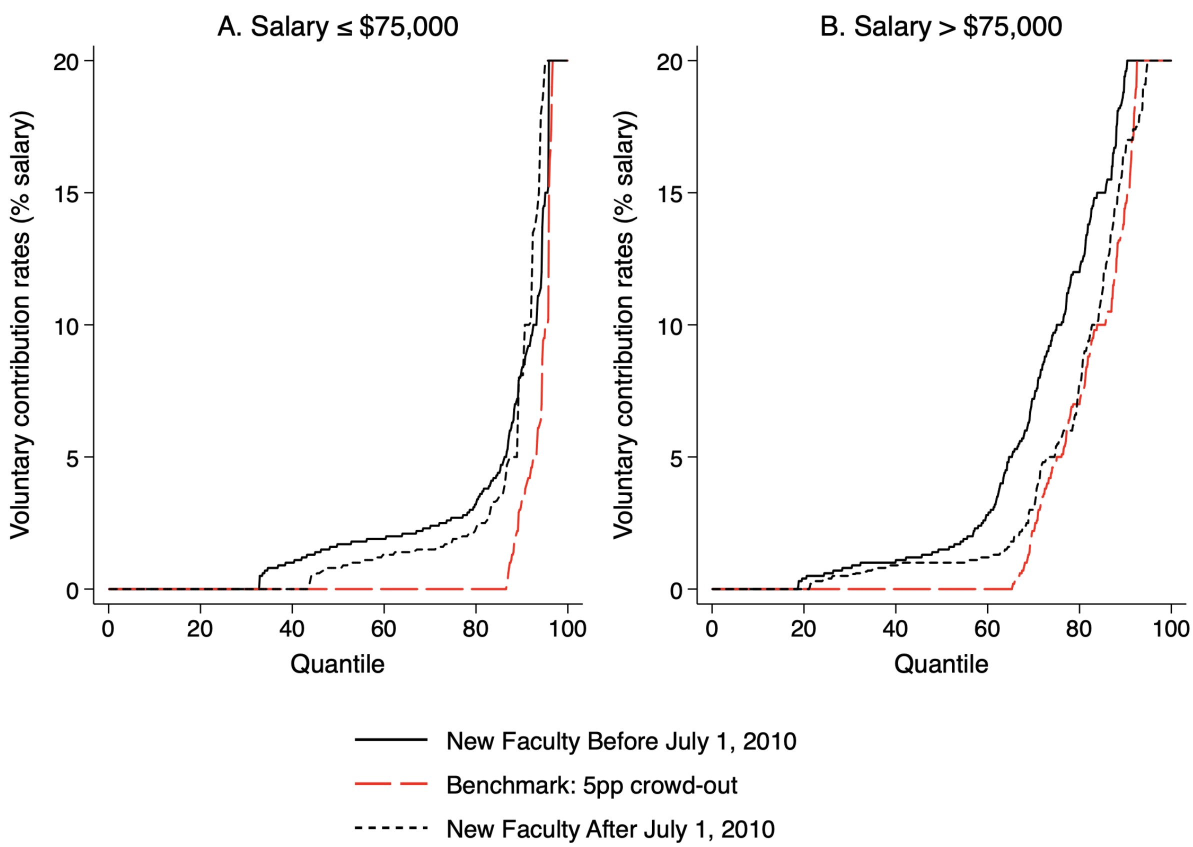

Notes: Figure plots the ordered values of voluntary contribution rates in the 9th month of tenure against quantiles, separately for faculty hired before July 1, 2010 and faculty and faculty hired after this date, in their 12th month of tenure. Panel (a) includes faculty with annual salaries less than $75,000 and Panel (b) includes faculty with salaries greater than or equal to this level. Sample excludes faculty who choose the DB plan. Voluntary DC contributions consist of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each; it excludes contributions that are smaller than 0.5% of salary, which represents the average contribution rate for a contribution of $40 per month, yielding the maximum allowable cash match of $20 per month. Roth contributions are adjusted for their post-tax status, and is reported relative to annual salary.

Online Appendices [Not for Publication]

A Supplementary Analyses [Not for Publication]

Figure A.1: Means of Demographics and Salary by Year of Hire, 2005-2014

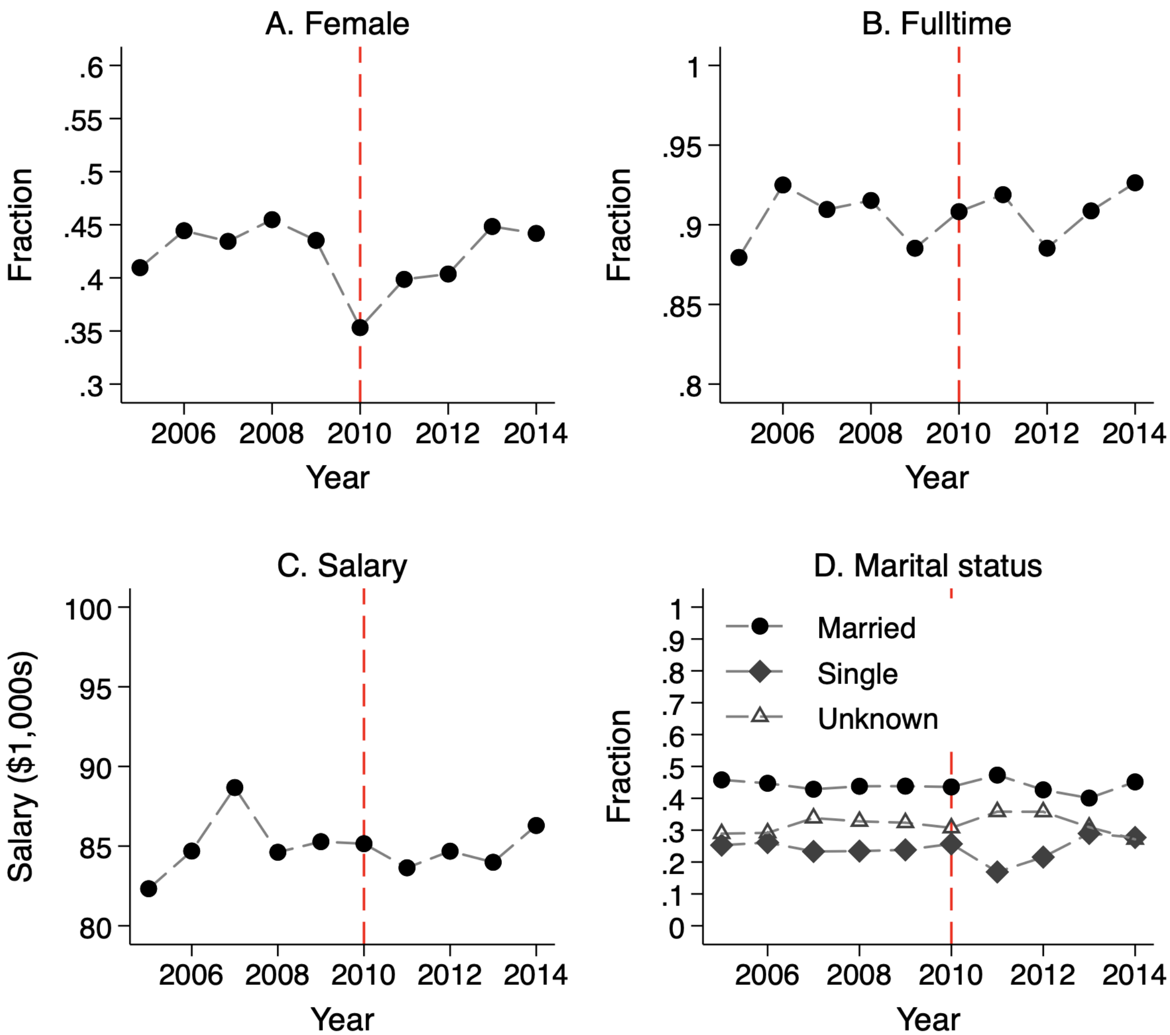

Notes: Administrative data on faculty at a large public university, statistics from their first month of hire. Figures plot the averages among employees hired each year from 2005 to 2014. The vertical line at 2010 denotes the year of the policy change.

Table A.1: Mean Participation and Contribution Rates in U.S. from Survey of Consumer Finances

	Participation (%)	Contribution (% salary)
2004	81.8%	8.3%
2007	79.0%	7.9%
2010	82.5%	8.0%
2013	84.0%	7.8%
2016	81.6%	7.9%

Table reports means weighted using survey sample weights from the Survey of Consumer Finances from the 2004, 2007, 2010, 2013, and 2016 surveys. We restrict analysis to employees whose employers offer DC plans to maintain comparability with our setting. We further restrict to employees who have a BA degree or higher, and who earn annual salaries exceeding $50,000. The means are calculated using the survey's sample weights and correspond to the individual level (the respondent answers for the spouse, if applicable, whose information we also include if they are also offered a DC plan).

Figure A.2: Quantile regressions of voluntary contribution rates

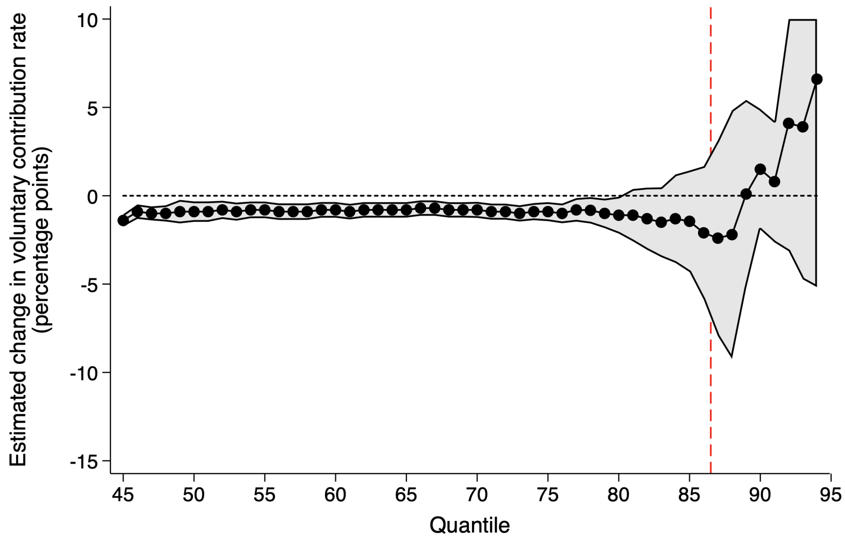

(a) Salary < $75,000

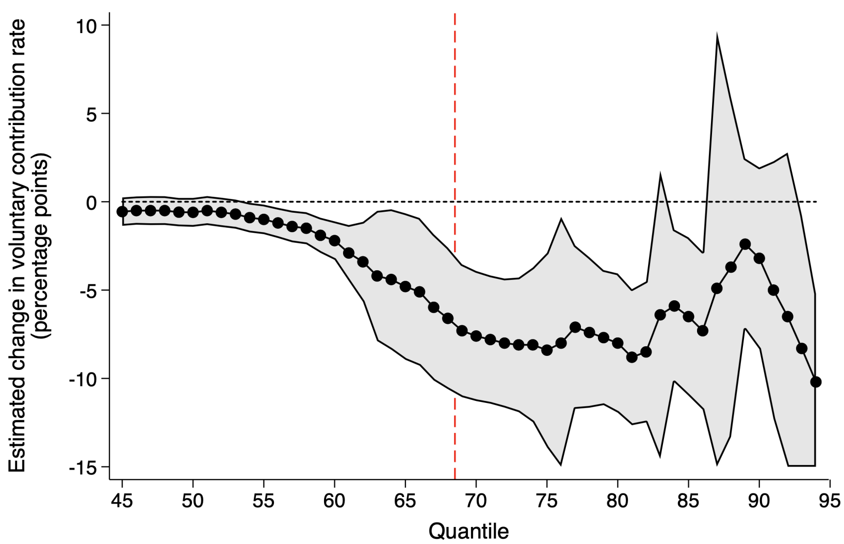

(b) Salary ≥ $75,000

Notes: Figure plots coefficient estimates from unconditional quantile regressions using the method of Firpo (2007), separately for faculty with salaries less than $75,000 (panel A) and $75,000 or above (panel B). Shaded regions represent 95 percent confidence intervals calculated using robust standard errors. For visual clarity, we do not display coefficient estimates below the 45th percentile since they are mechanically zero or close to zero, and we do not show coefficient estimates above the 95th percentile because the estimates are sometimes noisy and increase the y-scale necessary to present the results. The vertical line in each panel denotes the level of voluntary contributions corresponding to 5 percent of salary.

Figure A.3: Robustness: Voluntary participation and contribution rates by tenure month without controls, newly hired faculty

(a) Voluntary participation rate

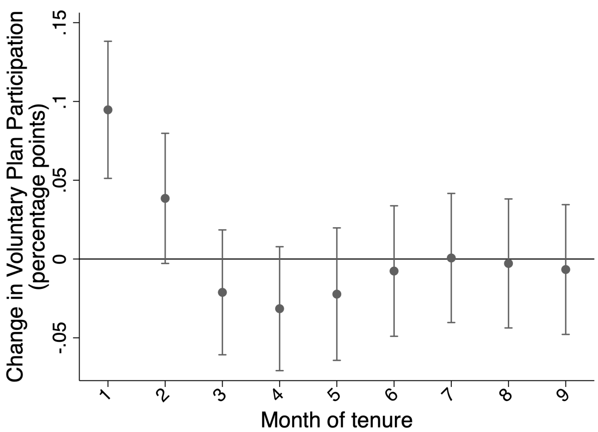

(b) Voluntary contribution rate, if positive

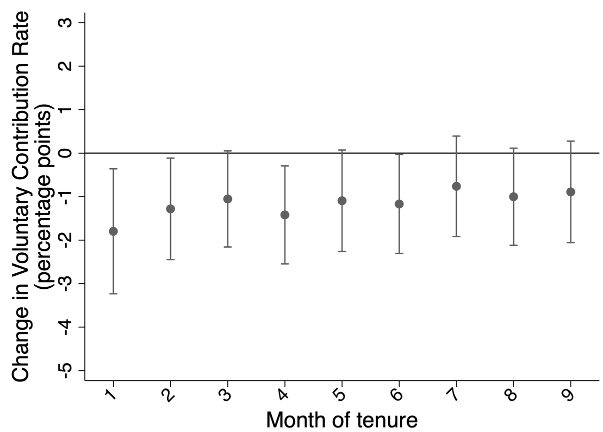

(c) Voluntary contribution rate

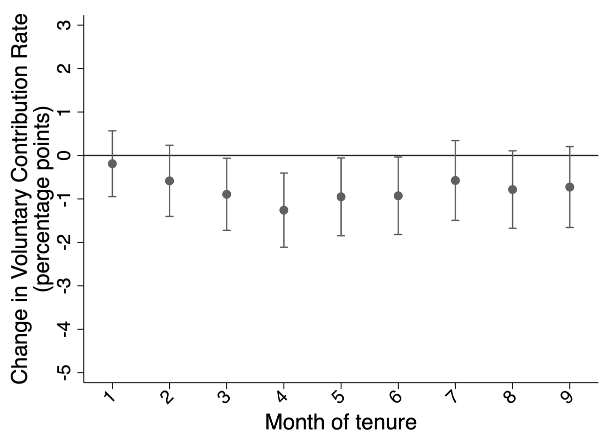

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation 4. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. Additional include indicators for tenure month and a constant but exclude other controls. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty do not receive a salary. Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each. The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary.

Figure A.4: Robustness: Voluntary participation and contribution rates by tenure month, newly hired faculty, excluding Sept 2008 – Dec. 2009

(a) Voluntary participation rate

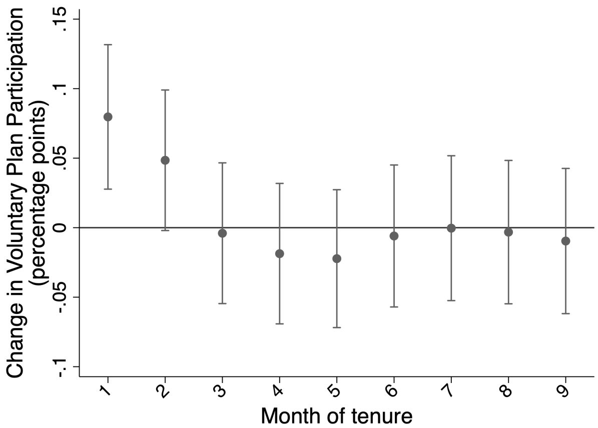

(b) Voluntary contribution rate, if positive

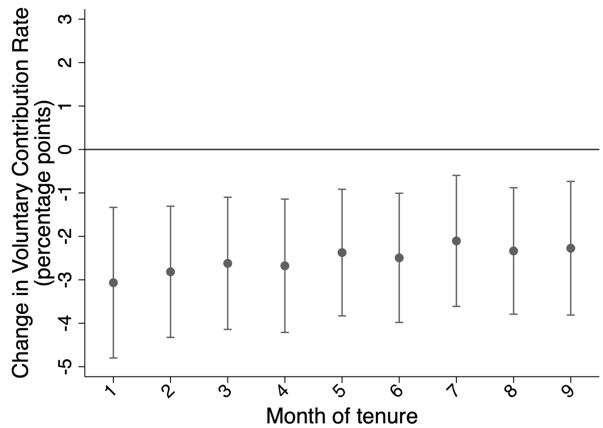

(c) Voluntary contribution rate

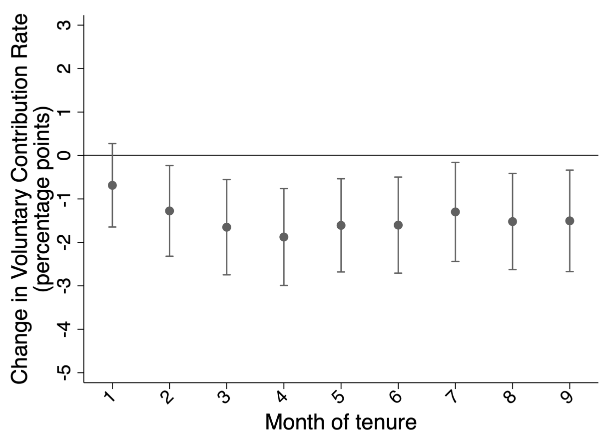

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation 4. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. The sample excludes observations from September 2008 to December 2009, inclusive. Additional control variables include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty do not receive a salary. Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each. The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary.

Figure A.5: Voluntary participation and contribution rates by tenure month, newly hired faculty, first 3 years

(a) Voluntary participation rate

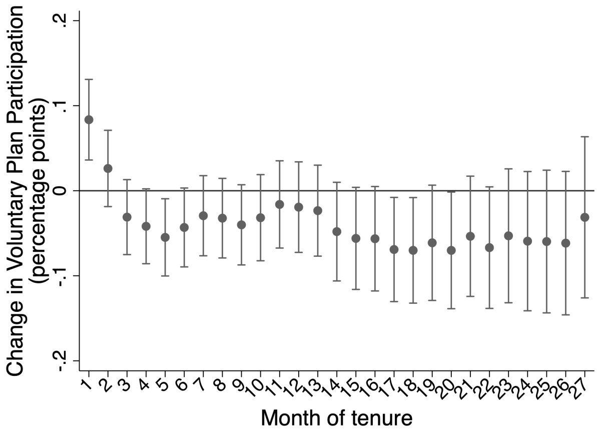

(b) Voluntary contribution rate, if positive

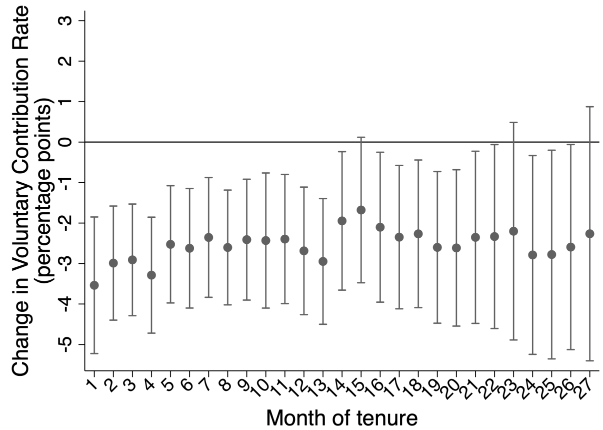

(c) Voluntary contribution rate

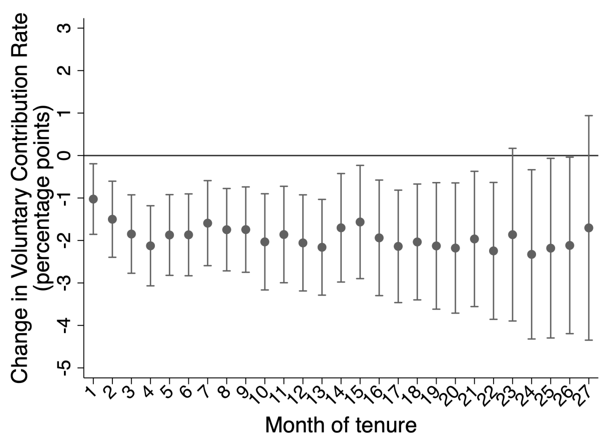

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation 4. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. Additional control variables include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands. Administrative data on faculty at a large public university, statistics from their first twenty-seven months of hire, which covers their first three years excluding June through August, when many faculty do not receive a salary. Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each. The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary.

Figure A.6: Placebo Test: Choice of Vendor in Retirement Accounts

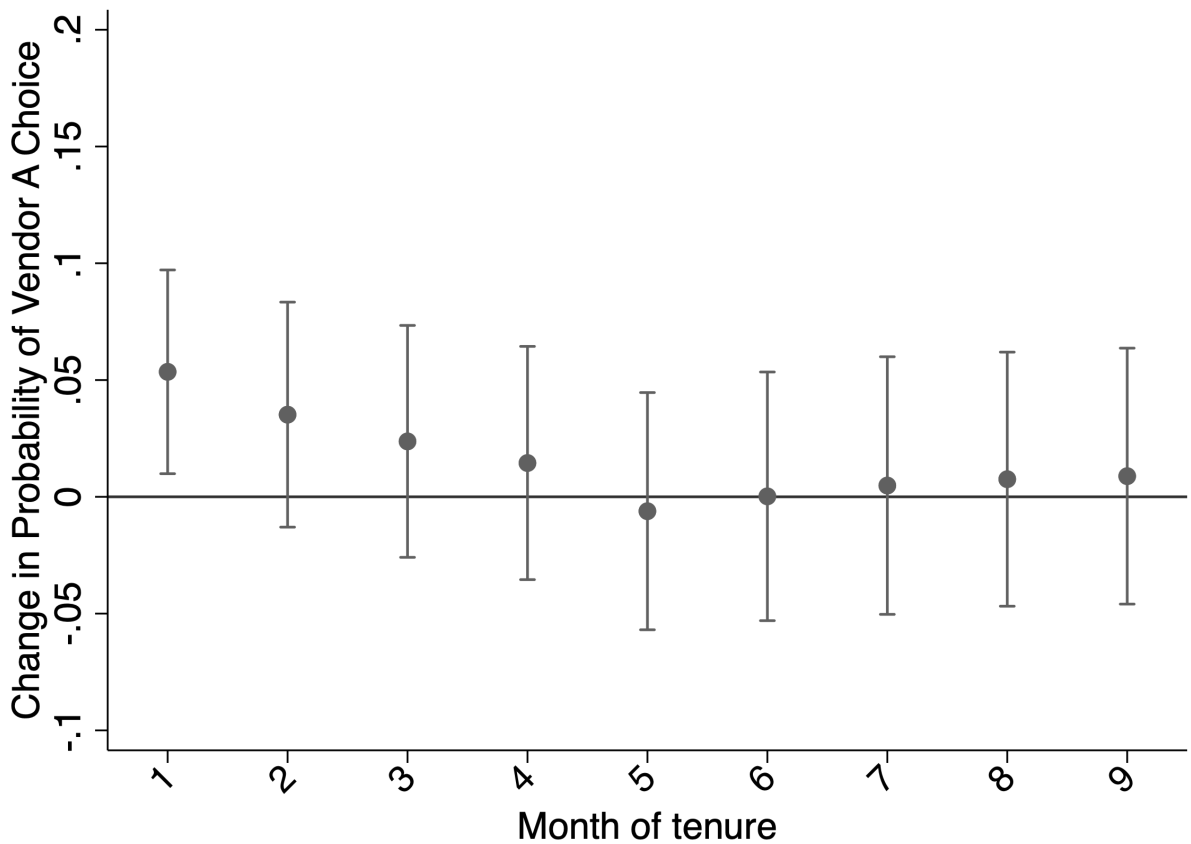

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation (3). The dependent variable is an indicator for choosing Vendor A (instead of Vendor B) as the account for voluntary contributions. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. Additional control variables include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty do not receive a salary.

Figure A.7: Regression-Adjusted Total Contribution Rates by Faculty/Staff and Year of Hire

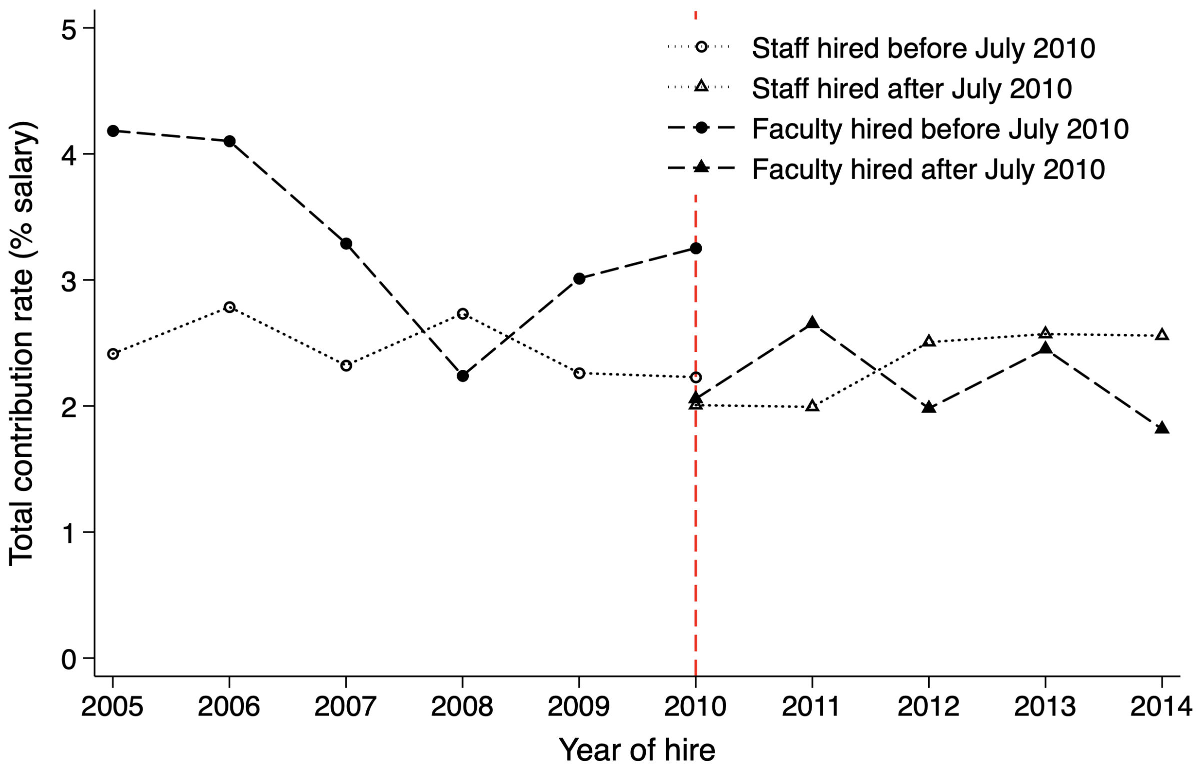

Notes: Figure plots mean total voluntary contribution rates for four groups: faculty hired before the policy change (shaded circles), faculty hired after the policy change (shaded triangles), staff hired before the policy change (hollow circles), and staff hired after the policy change (hollow triangles). Contribution rates average months 3 through 9 of the first year of tenure, and are regression-adjusted for age, income, gender, full-time status, and the level of the S&P 500 on the first of the month. There is little change over time in contribution rates by staff, who are unaffected by the policy. Among faculty, contribution rates are lower for employees hired after the policy change. All faculty are included in the graph, including those who choose the DB plan.

Figure A.8: Probability Contribute at Maximum

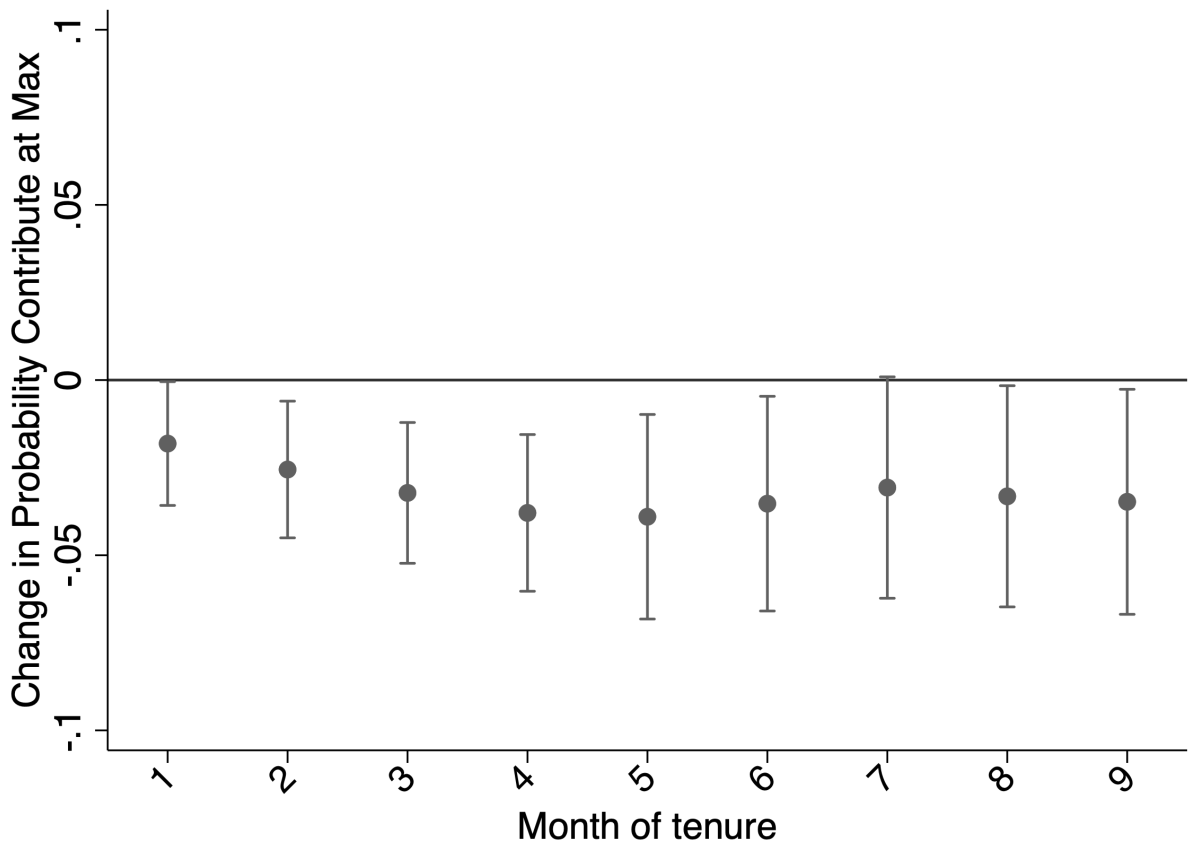

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation (3). The dependent variable is an indicator for making the maximum contribution to the voluntary plan. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. Additional control variables include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty do not receive a salary.

Table A.2: OLS Regression Estimates, DB Plan Choice

	Estimate	SE
Post-2010 policy change	0.023	(0.021)
Age, relative to $[25, 30)$		
$[30, 35)$	0.024	(0.0303)
$[35, 40)$	0.048	(0.0314)
$[40, 45)$	0.054	(0.0340)
$[45, 50)$	0.053	(0.0387)
$[50, 55)$	0.043	(0.0418)
$[55, 60)$	0.074	(0.048)
[60+)	0.111	(0.050)
Income, relative to $[20k - 25k)$		
$[25k - 29k)$	-0.082	(0.045)
$[30k - 35k)$	-0.009	(0.049)
$[35k - 39k)$	0.080	(0.051)
$[40k - 45k)$	0.265	(0.052)
$[45k - 49k)$	0.217	(0.050)
$[50k - 55k)$	0.119	(0.049)
$[55k - 59k)$	0.019	(0.054)
$[60k - 65k)$	-0.001	(0.052)
$[65k - 69k)$	0.009	(0.053)
$[70k - 75k)$	-0.029	(0.053)
$[75k - 79k)$	-0.031	(0.058)
$[80k - 85k)$	-0.115	(0.053)
$[85k - 89k)$	-0.108	(0.053)
$[90k - 95k)$	-0.139	(0.050)
$[95k - 99k)$	-0.133	(0.052)
$[100k - 120k)$	-0.192	(0.038)
$[120k - 140k)$	-0.075	(0.053)
$[140k+)$	-0.137	(0.044)
N	2756	

Table plots regression results of linear probability models of DB plan choice using administrative data on faculty at a large public university. Coefficient estimates with robust standard errors in parentheses. Faculty face a one-time irrevocable choice when hired between the state DB plan and the university DC plan. Age and annual salary are reported in bins to preserve confidentiality of individuals in the data set. Regressions also include indicators for sex, fulltime employee, marital status, the level of the S&P 500 on the 1st of each month, and a constant.

Figure A.9: Robustness: Voluntary participation and contribution rates by tenure month, all newly hired faculty

(a) Voluntary participation rate

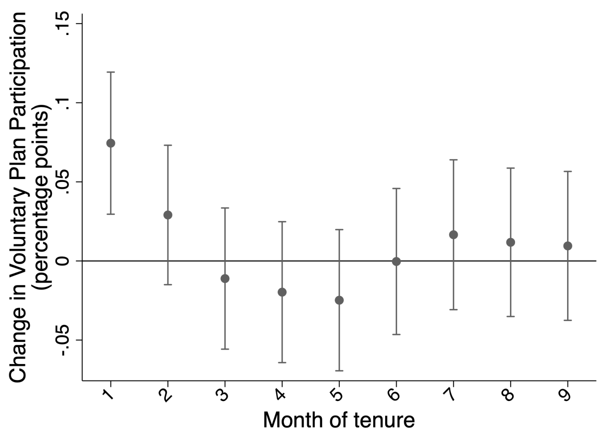

(b) Voluntary contribution rate, if positive

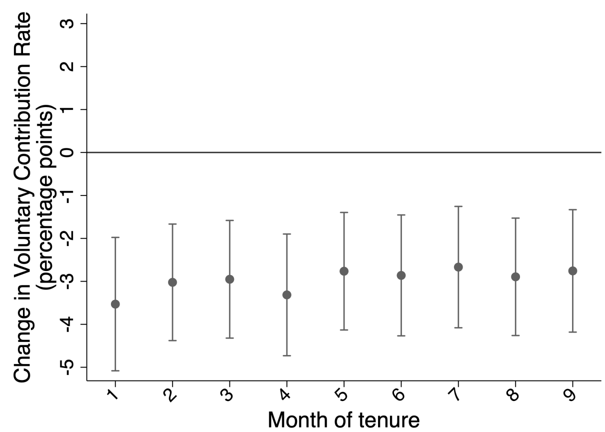

(c) Voluntary contribution rate

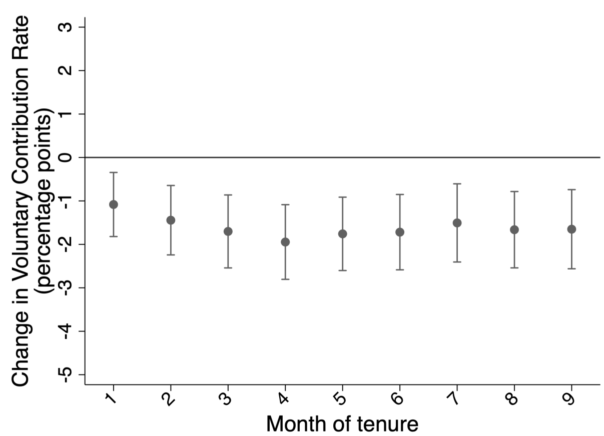

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation 4. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. The sample includes all newly hired faculty, including those who choose DB plan. Additional control variables include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty do not receive a salary. Faculty who chose the university DC plan with mandatory contributions when hired can then choose voluntary DC contributions. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each. The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary.

Table A.3: Descriptive Statistics by Salary Level

	Pre: Hired 2005-2009		Post: Hired 2010-2014		Diff.	p-value
	Mean	SD	Mean	SD		
Panel A. Salary below $75,000						
Voluntary DC participation rate	0.683	0.466	0.643	0.480	-0.039	0.279
Voluntary DC contribution rate, %	3.449	7.341	3.502	8.842	0.053	0.930
Voluntary DC contribution rate, % (if pos.)	5.555	11.251	6.156	13.718	0.601	0.598
Demographic and job characteristics:						
Female	0.491	0.500	0.500	0.501	0.009	0.819
Married	0.415	0.493	0.404	0.492	-0.011	0.770
Single	0.311	0.463	0.322	0.468	0.011	0.760
Marital status unknown	0.274	0.446	0.274	0.447	0.000	0.996
Full-time	0.847	0.360	0.835	0.372	-0.012	0.663
Annual salary ($1000s)	49.90	14.22	46.77	15.26	-3.13	0.005
Age	37.56	8.64	36.71	7.48	-0.85	0.187
S&P 500 Index	1263	209	1614	252	351	0.000
p-value from omnibus balance test						0.191
Panel B Salary above $75,000						
Voluntary DC participation rate	0.821	0.384	0.807	0.395	-0.014	0.511
Voluntary DC contribution rate, %	6.600	9.859	4.799	8.484	-1.801	0.000
Voluntary DC contribution rate, % (if pos.)	8.372	11.944	6.044	9.547	-2.328	0.000
Demographic and job characteristics:						
Female	0.374	0.484	0.402	0.491	0.028	0.289
Married	0.465	0.499	0.442	0.497	-0.023	0.386
Single	0.176	0.381	0.213	0.410	0.038	0.076
Marital status unknown	0.359	0.480	0.345	0.476	-0.014	0.578
Full-time	0.957	0.204	0.955	0.208	-0.002	0.853
Annual salary ($1000s)	122.82	32.90	121.96	29.78	-0.86	0.676
Age	41.29	9.09	39.68	9.37	-1.61	0.001
S&P 500 Index	1259	220	1777	287	517	0.000
p-value from omnibus balance test						0.881

Notes: Descriptive statistics calculated separately on sub-samples by salary level. Panel A presents statistics for employees with 1st year salaries below $75,000, and Panel B presents statistics for employees with 1st-year salaries greater or equal to $75,000. The bottom row of each panel presents the p-value from the omnibus balance test of demographics and job characteristics, excluding the S&P index.

Table A.4: Composition of employees staying after 18 and 27 months

	Pre: Hired 2005-2009 Mean	Post: Hired 2010-2014 Mean	Difference	p-value
If still employed after 18 months				
Female	0.423	0.451	0.029	0.394
Married	0.484	0.403	-0.081	0.018
Single	0.211	0.257	0.046	0.104
Marital status unknown	0.305	0.340	0.035	0.272
Full-time	0.929	0.910	-0.019	0.296
Annual salary	99.614	103.902	4.288	0.157
Age	41.239	40.661	-0.578	0.348
p-value of omnibus balance test				0.059
If still employed after 24 months				
Female	0.410	0.367	-0.044	0.512
Married	0.492	0.433	-0.059	0.386
Single	0.205	0.233	0.028	0.610
Marital status unknown	0.303	0.333	0.031	0.624
Full-time	0.948	0.917	-0.031	0.315
Annual salary	102.681	107.614	4.934	0.417
Age	41.204	39.971	-1.234	0.282
p-value of omnibus balance test				0.537

Table plots mean characteristics for employees who are still employed after 18 months (2 years, excluding summer months) and after 27 months (3 years, excluding summer months), separately by whether the employee was hired prior to 2010 or after 2010. The third column displays the difference in means, and the fourth column reports the p-value from the t-test that this difference is equal to zero. In several cases, the differences are statistically different, reflecting that that the composition of those who remain employed is systematically different for employees hired before 2010 versus afterwards.

Table A.5: Descriptive Statistics of Staff

	Pre: Hired 2005-2009		Post: Hired 2010-2014		Diff.	p-val
	Mean	SD	Mean	SD		
Voluntary DC participation rate	0.460	0.498	0.539	0.499	0.079	0.000
Voluntary DC contribution rate, %	1.782	3.639	1.810	3.610	0.028	0.763
Voluntary DC contribution rate, % (if positive)	3.897	4.552	3.367	4.359	-0.530	0.001
Demographic and job characteristics:						
Female	0.537	0.499	0.515	0.500	-0.022	0.088
Married	0.161	0.367	0.107	0.309	-0.054	0.000
Single	0.202	0.401	0.191	0.393	-0.011	0.269
Marital status unknown	0.637	0.481	0.703	0.457	0.065	0.000
Full-time	0.953	0.211	0.955	0.207	0.002	0.705
Annual salary	43.9	15.3	45.7	17.2	1.8	0.000
Age	36.7	9.9	36.5	10.2	-0.2	0.420
S&P 500 Index	1242	208	1649	299	407	0.000
p-value from omnibus balance test						0.000
N	3632		2701			

Administrative data on staff at a large public university, statistics from their ninth month of hire. The final column lists the p-value from the t-test the means are equal between employees hired in 2005-2009 vs. employees hired in 2010-2014. Voluntary DC participation consists of contributions to either the 403(b) plan or the 457 plan, or, beginning in 2013, Roth options of each. The voluntary DC contribution rate is defined similarly, with Roth contributions adjusted for their post-tax status, and is reported relative to annual salary. Age and annual salary are reported in bins to preserve confidentiality of individuals in the data set. Age is grouped into bands, generally of five years, and we impute individual age as the mid-point of the age bands. We impute salary as the midpoint of salary bands, winsorized at the 5th and 95th percentiles to remove the influence of outliers on the mean. Both salary and the S&P index are inflation-adjusted to 2018.

Table A.6: Difference-in-Differences Regressions of Total Contribution Rates

	(1)	(2)	(3)	(4)
Faculty x Post-2010	-0.701	-0.883	-1.014	-1.067
	(0.368)	(0.369)	(0.329)	(0.327)
Faculty (1=yes, 0=no)	3.249	1.420	3.034	1.220
	(0.232)	(0.230)	(0.217)	(0.216)
Post-2010 (1=yes, 0=no)	-0.049	-0.155	0.025	-0.119
	(0.346)	(0.346)	(0.339)	(0.337)
Sample	Staff + faculty in DC plan	Staff + faculty in DC plan	Staff + all faculty	Staff + all faculty
Controls	No	Yes	No	Yes
Observations	52888	51665	54983	53740
R^2	0.056	0.113	0.048	0.108

Table plots difference-in-difference regressions of total voluntary contribution rates, comparing new faculty to new staff hired before and after the July 1, 2010 policy change. Regressions include the level of the S&P 500 on the first of the calendar month, and indicators for age, income, gender, fulltime status, tenure month, and year of hire. The regression includes rates from tenure months 3-9, inclusive, since employee contribution rates are often low in the initial quarter after being hired.

Table A.7: Welfare Analysis among Employees with Salaries Below $75,000: Consumption Equivalents for Active Low Savers and Passive Savers

	Lower Bound on Passive Savers	Upper Bound on Passive Savers
Panel A. Percentage of each type		
Passive savers	56	89
Active low savers	44	11
Active high savers	0	0
Panel B. Costs and Benefits ($/year), CRRA = 3		
Cost: Avg. consumption equivalent, active low savers	112	94
Benefit: Avg. consumption equivalent, passive savers	223	213
Ratio of benefits to costs, weighted by shares in (A)	2.54	18.31
Panel C. Costs and Benefits ($/year), CRRA = 4		
Cost: Avg. consumption equivalent, active low savers	158	135
Benefit: Avg. consumption equivalent, passive savers	284	271
Ratio of benefits to costs, weighted by shares in (A)	2.29	16.26

Notes: Table summarizes results of consumption equivalent calculations for different types of savers among employees with annual salaries of $75,000 or less. Panel A presents the percentage of savers classified according to the types based on the quantile plots from Figure 5. The first column considers the lower bound on the fraction of the sample who are passive savers, which is equivalently the upper bound on who are active low savers. The second column considers the upper bound on passive savers, equivalently the lower bound on active low savers. Panels B and C show the average certainty equivalents for active low savers (costs) and passive savers (benefits) for CRRA coefficients of 3 and 4, respectively. The ratio of benefits to costs is calculated as the average benefits weighted by the percentages in Panel A of each type.

B Imputation of Marginal Tax Rates [Not for Publication]

This Appendix describes the procedure to impute marginal tax rates for each employee in our data. The marginal tax rates are used to adjust contributions to Roth accounts, which became available in in the later period we examine, to a pre-tax basis. Our administrative records lack several pieces of information required for a direct calculation of the employee's marginal tax rate, including information about spousal earnings, children, other sources of income, home ownership, and relevant deductions. In addition, marital status is reported incompletely and salary is recorded in bands to protect data confidentiality. Our approach is therefore to calculate marginal tax rates for respondents of the American Community Survey (ACS) using the National Bureau of Economic Research's TAXSIM, and then to use hot-deck imputation to assign a marginal tax rate for the employees in our sample by matching on income, age, and gender.

Step 1: ACS data We use ACS surveys between 2011 and 2017, which record relatively comprehensive information that helps us calculate marginal tax rates. In particular, we use the following information from the survey: wage and salary income of respondent and spouse, interest received, retirement income and social security benefits, supplemental security income and public assistance income, state, marital status, age, number of dependents, and number of children under 13.

Step 2: Marginal tax rate calculation For each ACS observation, we use NBER TAXSIM to estimate the federal and state marginal tax rates based on the variables in the list above.

Step 3: Hot-deck imputation We match individuals between our administrative data and the ACS by year, age band, income band, and gender. We then use hot-deck imputation to assign a marginal tax rate to the matched employees in our sample. The imputation is repeated five times and we take the average to construct our estimate of the employee's marginal tax rate.