Does Mandatory Retirement Saving Crowd Out Voluntary Retirement Saving?*

Leora Friedberg
University of Virginia

Adam Leive
University of California-Berkeley

Wenqiang Cai
PriceWaterhouseCoopers LLP

October 2023

Abstract

We use administrative data from a large public university to study employee responses to mandatory retirement saving, which represent a policy alternative to nudges. In response to a substantial increase in mandatory saving, we observe full crowd-out of voluntary saving for a subset of low-contributing low earners and high-contributing high earners, but otherwise little crowd-out for most employees. On average, only 30 percent of the mandatory contributions are offset by lower voluntary contributions. Our results suggest that mandatory retirement saving is likely to increase total retirement saving.

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 NBER Summer Institute (Aging), UVA Applied Microeconomics Workshop, New School for Social Research, TIAA Fellows Institute Symposium on Behavioral Finance, and National Tax Association Meeting for helpful comments. Alex Sheng and 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 who have an employer pension today have a defined contribution (DC) plan.[1] A key feature of DC plan design is the considerable autonomy that rests with individuals over how much to contribute, giving them major responsibility 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 the choice to opt out (Thaler and Sunstein 2008). Extensive research demonstrates the power of defaults in increasing short-run participation in DC plans (Madrian and Shea 2001, Thaler and Benartzi 2004), but recent evidence casts doubt on their effectiveness in raising savings rates in the long run (Choukhmane 2021, Laibson 2020). Related theoretical research highlights stronger paternalism mechanisms, notably mandatory contributions, that can raise not only long-term saving but also 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, which are a defining feature of public pension plans, therefore emerges from both recent evidence and theory.[2]

This stands in contrast to the neoclassical model, in which mandatory saving fully crowds out voluntary saving for rational individuals—except for individuals whose optimal saving is less than the mandatory savings rate and who cannot respond with full crowd-out. In this paper, we analyze changes in voluntary DC contributions in response to an increase in mandatory employee contributions. We analyze not only the mean crowd-out response but also the magnitude of crowd-out along different parts of the distribution of voluntary contributions. Then, based on the distribution of observed crowd-out, we apply a simple economic framework that can guide an understanding of consumption-equivalent values, which indicate by how much an employee must be compensated, or alternately would be willing to pay, under differing assumptions to make mandatory contributions.

We use administrative data from a public university that instituted an increase in mandatory contributions that was large and plausibly exogenous. The contribution rates are close to what is observed in the public setting, and are notably greater than those studied in other recent private settings (for example, in Chetty et al. 2014). The policy change followed state legislation that focused on funding the state defined benefit (DB)

[1] Since the decline of 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 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). Defaults would not solve this failure.

plan for all state employees, so we view it as exogenous for new employees in the university DC plan, rather than, for example, an endogenous response to employee preferences. Starting in 2010, newly hired faculty experienced two changes relative to earlier hires: 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 (a 5 pp increase, effectively). Consequently, the total mandatory contribution rate rose by 3.5 pp. As predicted by standard economic theory, 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 on whether salaries adjust or not for the change in deferred compensation and whether there is a resulting income effect), and a decline to zero at the low end, as low savers cannot fully crowd out the new mandatory contribution.

Our primary research design has two components: it compares the full distribution of voluntary DC contributions and not just the mean, and it compares those distributions for employees hired after the increase in mandatory contributions to employees hired before. The comparison of the full distribution of voluntary contributions sheds light on the source of crowd-out and on active versus passive responses. The comparison between cohorts of new employees follows much of the literature on defaults and other retirement savings policies (Madrian and Shea 2001, Carroll et al. 2009, Beshears et al. 2022).[3] We implement an alternative research design, finding quantitatively similar estimates using new staff—who were not directly affected by the policy change—as a control group for new faculty in a difference-in-differences design. This helps address concerns about employer-specific changes in hiring or information provision or economy-wide changes in savings choices. Also, using the Survey of Consumer Finances, we demonstrate that DC plan contribution rates did not change observably for a similar sample of nationally representative employees.

We find strong evidence that crowd-out is incomplete when considering mean responses. Voluntary contribution rates overall fell by about 1.5 pp, and we reject the full crowd-out benchmark of 3.5–5 pp.[4] Participation in the voluntary account fell slightly, but not by enough to explain incomplete crowd-out.

Recognizing that deviations from full crowd-out may reflect various underlying responses (for example, partial crowd-out by everyone or full-crowdout by some), we then

[3]This approach illustrates the theoretical rationale offered in Fadlon and Laibson (2019) for leveraging policy changes to evaluate whether households optimize. They show that empirical tests based simply on cross-sectional savings or consumption data cannot distinguish responses to mechanisms imposed by social planners that correct optimization failures, from optimal individual behavior.

[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 salaries rose insignificantly, and we cannot reject a fully-offsetting rise of 1.5 pp, which would imply full crowd-out of 5 pp.

investigate the distribution of voluntary contributions to consider where crowd-out originates. We further split the sample at median earnings (of about $75,000) to reveal the types of employees who respond, since lower earners may save less and/or face more liquidity constraints. The distributional analysis demonstrates four types of responses:

1. *Active high savers*: employees saving more than the mandated increase who respond with full crowd-out (comprising 17% of the full sample).

2. *Active low savers*: employees saving less than the mandated increase who stop making voluntary contributions (8%).

3. *Passive savers*: employees who display little or no crowd-out and continue making voluntary contributions (50%).

4. *Non-savers*: employees not making voluntary contributions in the first place (25%).

Moreover, the full crowd-out responses that we observe occur 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 parts of these savings distributions, we observe very little responsiveness by many employees (passive savers). The incomplete crowd-out that we observe thus largely reflects a relatively small share of employees who respond fully.

We finish our analysis with calculations of consumption-equivalent values under different scenarios, using information from the responses that we observe. These are the dollar amounts that a person would be willing to forgo each period to avoid making mandatory contributions while working—or alternately, that they would be willing to pay to be required to make them. While we know the magnitude of active versus passive responses, we do not know whether someone saving passively, nor even actively, is nudged towards or away from their optimal savings rate. Instead, we consider some extreme cases. For example, if everyone were saving optimally before mandatory contributions were raised, then the policy induces too much saving by active low savers, passive savers, and non-savers, tightening their liquidity constraints. In this case, the consumption equivalent averaged over the full sample is $1,765 (about a 3.0% loss in consumption). Another extreme possibility is that the policy moves everyone except active savers with full crowd-out to their optimal saving rate: in this case, the consumption equivalent is about negative $2,225 (about a 4.0% gain in consumption). In reality, some employees were likely under-saving without the policy while others were already saving optimally, or even over-saving. One qualitative takeaway from this analysis is that the consumption equivalents are large in magnitude for lower earners because they exhibit less crowd-out.

Given that we do not know whether they are positive or negative, future research will need to explore the balance between liquidity and commitment for this group.

Our results contribute to research on private and public 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.[5] The two most closely related papers both document passive saving behavior in response to mandated private-pension contributions.[6] Chetty et al. (2014) study the impact of a 1 pp increase in mandatory contributions in Denmark and find a negligible reduction in pension saving. Notably, we find a similarly small percentage of active responses that exhibit full crowd-out, even though the mandatory increase in our setting is five times larger. Thus, our findings appear inconsistent with rational inattention, as the larger change in mandatory contributions in our setting should elicit a greater response.[7] In other related work, Card and Ransom (2011) study crowd-out of voluntary contributions among faculty across universities in the 1980s and 1990s. They estimate crowd-out rates of between 60 and 80% of the theoretical benchmark in response to mandated employee contributions, but rates only half as large in response to employer contributions. While we cannot separately distinguish the response to the employee versus employer-mandated contribution in our setting, our results provide further insight 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%, may reflect an increasingly complicated choice environment that employees face across their workplace benefits (Leive et al. 2022).

The rest of this paper is organized as follows. Section 2 describes the institutional details of the large public university that we study. Section 3 introduces a theoretical model that predicts how average contributions to voluntary plans change when total contributions to the mandatory plan increases. Sections 4 to 6 describe, respectively, the data, empirical strategy, and the main results of this paper. Section 7 concludes.

[5] Bubb and Warren (2020) differs in studying employer rather than social-planner motives for paternalistic pension design. Their model suggests further limitations of nudges.

[6] A large literature studies crowd-out across a range of public and private pension settings (Feldstein 1974, Munnell 1976, Hubbard and Judd 1986, Hubbard 1998, Bottazzi and Padula 2006, Aguila 2011, Attanasio and Rohwedder 2003, Attanasio and Brugiavini 2003, Engelhardt and Kumar 2011, Hurd and Rohwedder 2012, Lachowska and Myck 2018). Those estimates suggest that a \$1 of mandated saving crowds out anywhere between roughly \$0.25 to \$0.75 of private saving. Yang (2020) finds that when a Taiwanese law required employer contributions of 6% of salary to a retirement account, employees reduced private saving by roughly 2%-2.5%. Many of those studies are relevant in demonstrating that total savings and not just other workplace accounts exhibit far less than full crowd-out.

[7] The policy analyzed by Chetty et al. (2014) imposed a required contribution of 1% of earnings for some Danish citizens. The study also analyzed variation in employer contribution rates (equal to 5% of earnings) among employees who switch firms, and found a quite similar rate of crowd-out.

2 Institutional Details

The large public university that we study offers faculty a complicated but relatively typical 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.[8] As we show in Section 4, about 80% choose the DC plan, a fraction that rose by a small and insignificant amount, in spite of the new mandatory employee contribution in the DC plan. Second, faculty must choose, for their mandatory contributions, among two vendors and numerous funds for each vendor.[9] Third, faculty can choose additional voluntary contributions, which are subject to a very limited match rate, of 50% for up to $40 per month and 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.[10] Fifth, faculty must choose among two vendors for 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 option that employees have in saving for retirement due to its tax preferences, low expenses, and ease of taking loans. The advantages relative to IRAs or taxable accounts is important since we only observe employer accounts. We therefore assume, as in much of the literature on employer pension plans, that the any change in contributions to these accounts represents a change in overall savings.

Conditional on choosing the university DC plan in the first step above, the required contribution rates in the plan depend 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

[8]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 occurred in the last year of our sample and did not noticeably change the DB plan choice.

[9]Employees can split contributions among the vendors and among all the funds that each vendor offers.

[10]Just like with 401(k) plans, employees can split their contributions to the 403(b) and Roth 403(b) plans and get tax preferences up to IRS contribution limits. Employees can additionally split their contributions to the 457 and Roth 457 options up to the same 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.

Figure 1: Flow Chart of Retirement Saving Choices

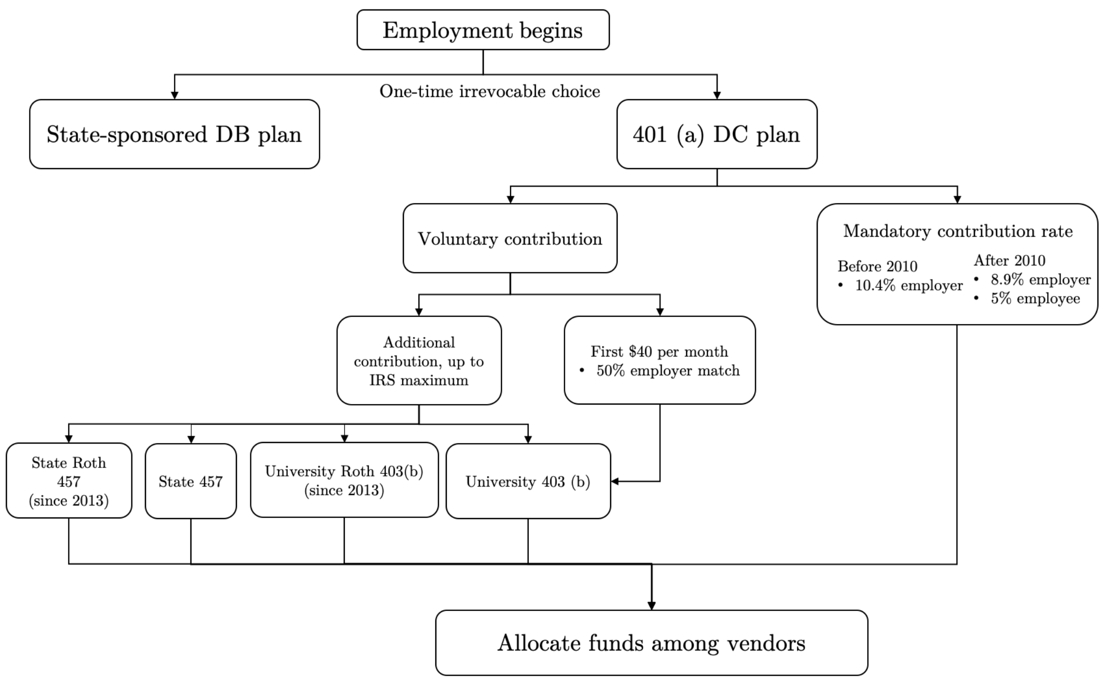

Notes: Figure displays the sequence of choices and available options for retirement plans at the university. The first choice is between the state-sponsored DB plan and the 401(a) DC plan. The main policy of interest is the change in the mandatory contribution rate in the 401(a) DC plan that occurred 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.

state DB plan; it also mandated a new 5% employee contribution for new hires among those few state employees (mostly faculty at a handful of state universities) who could opt for a DC plan, putting them on par with other state employees who were already required to contribute 5% to the DB plan.[11] 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 preferences of new faculty hires or an element of a broader set of changes in university policy.

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.

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. Two 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 at a contribution of $40, so that they would automatically benefit from the match. Then, in 2013, the Roth options mentioned above became available. Later on, we describe straightforward adjustments we make to the data to deal with these background policy changes.

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

[11]The changes in the DB plan involved raising the eligibility age for benefits, slightly reducing service credits, and reducing the cost-of-living adjustment. 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 lower contributions through a reduced-wealth effect; and, in any case, we observe almost no change in behavior.

3 Model

A stylized theoretical model of optimizing savers helps us understand conditions under which mandatory contributions would theoretically crowd out voluntary contributions. We will use this full crowd-out benchmark to evaluate whether observed crowd-out responses in our setting are consistent with optimizing behavior.[12] This analysis applies similarly to public pensions, which often involve some combination of required employee and employer contributions.

We highlight three predictions of the benchmark model for voluntary contributions, with associated welfare effects. One is a crowd-out effect on voluntary contributions, from the new required employee contribution. The second 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 third is a compensation effect from the reduction in the required employer contribution, to the extent that this is not compensated by an increase in salary. The compensation effect reflects a welfare loss for all employees, while the participation effect reflects a further welfare loss because some employees are forced to oversave because of higher mandatory contributions.

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, retiring in 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 savings. A strong savings motive in this baseline model means that an employee's savings in the absence of any required contributions exceeds the required contribution amounts that we consider.

Employees must contribute M_e to the mandatory retirement savings account, and the firm must contribute M_f. Employees then face a choice of contributing an additional non-negative amount M_v to a voluntary retirement savings account.[13] We assume that both

[12] 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, some optimizing agents respond by less because they have less savings to crowd out; in their case, full crowd-out means that their voluntary contributions drop to zero.

[13] 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.

mandatory and voluntary retirement saving accounts have the same rate of return r.

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:

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

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 a drop in the required employer contribution and a new required employee contribution, leading to a higher required total. We assume no change in income Y, so effective compensation falls.[14] She maximizes the same objective function as employee i but faces modified budget constraints:

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

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 in our setting. BL^i and BL^j represent the budget lines before and after the policy change, which has two effects. First it reduces total compensation, resulting in a parallel shift not of 5 pp but of 1.5 pp, resulting from the reduced employer contribution.

[14] If instead, salaries increase, leaving total employer compensation unchanged, then the compensation effect outlined below would not operate, leading to bigger predicted crowd-out. In fact, we observe a modest increase in imputed salary in our sample and cannot reject either a zero or full jump; but since our salary data is reported in ranges, we are not confident in making claims about the presence or absence of a compensation effect.

Second, it eliminates access to very high values of c_1 because of the new employee contribution of 5 pp.

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 distance between Y and c_1^i is M_v^i, the voluntary contribution amount for early hires. Similarly, the distance between $0.95Y$ and c_1^j is M_v^j, the contribution amount to the voluntary plan for later hires.

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. 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 Full Model with Non-Participation

The predictions in our baseline model pertain to the case of high voluntary contribution rates. If voluntary contributions for employees before the policy change are small or zero,

Figure 2: Graphical Representation of Theoretical Model

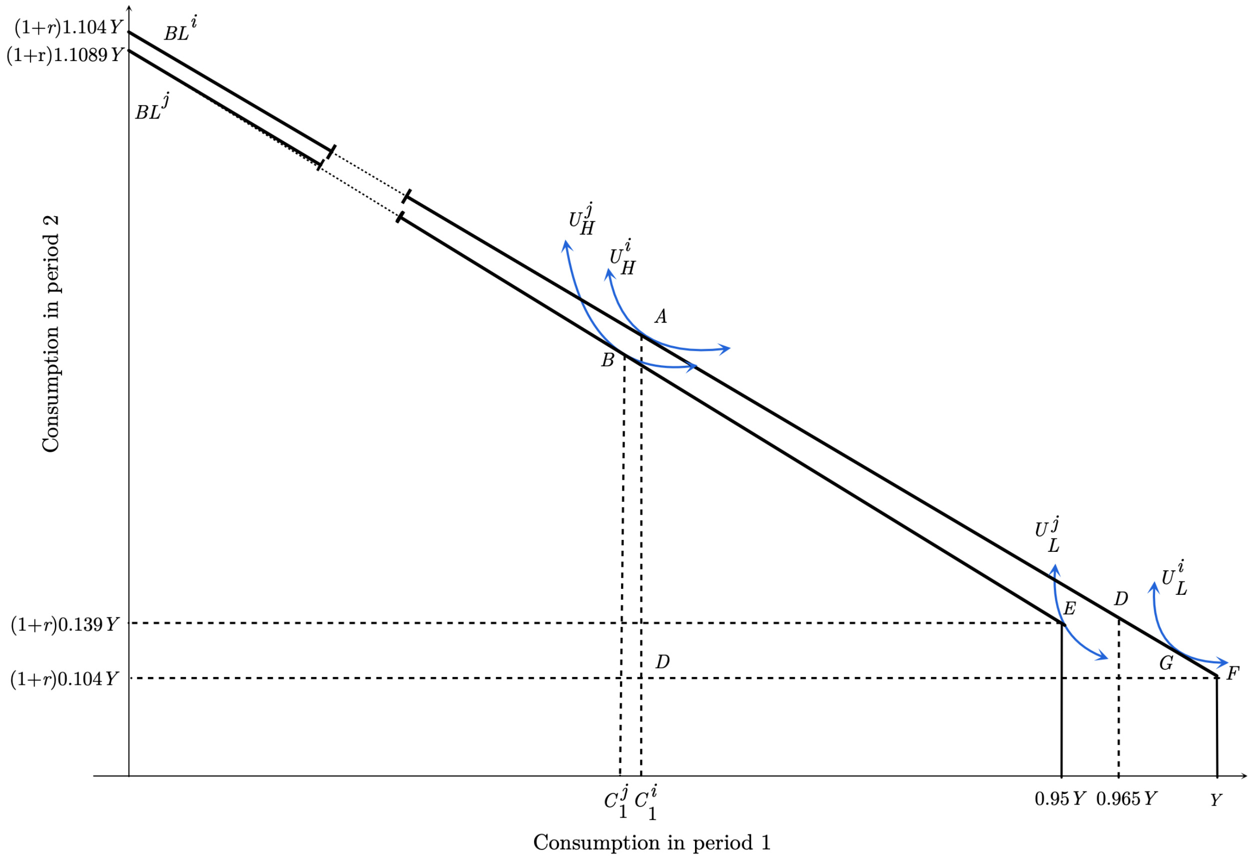

Notes: This graph sketches the predictions of the standard theory of retirement savings crowd-out. The budget line before the policy change (without mandatory contributions) is BL^i and the budget line after the change is shifted inwards to BL^j. Point A represents the optimal consumption bundle for 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. The dotted lines in the budget constraints represent extensions of each line so that the scale is easier to read.

however, that limits the crowd-out effect, replacing it instead with a participation effect.

In Figure 2, this case is illustrated by the 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.

Among such 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 smaller, because those saving above 5% salary can fully offset the mandatory increase through reduced voluntary contributions.

Table 2: Summary of Model Predictions

	Voluntary contribution rate	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 to zero	Stop participating

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. 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.

3.3 Discussion

Table 2 summarizes theoretical predictions for the full crowd-out benchmark with optimizing savers. These predictions apply to private or public pension plans. For individuals who would save a considerable amount counterfactually, the average voluntary contribution rate will decrease by 3.5 to 5 percentage points, depending on whether a compensation reduction partially offsets the crowd-out effect. For individuals 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 individuals 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 and younger workers are more likely to face binding liquidity constraints, we will split our sample by earnings and age when investigating the distributional responses.[15]

4 Data

To study responses to mandatory contributions, 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 demographic 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.[16] 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 financial conditions, we merge in monthly levels and changes 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

[15] Of course, those with high incomes may also face liquidity constraints (Kaplan et al. 2014.)

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

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 2,832 new faculty hired between 2005-2014. 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 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.[17]

We consider contribution rates of new faculty nine months after being hired, when it is apparent that participants reach a steady contribution rate.[18] 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, 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 versus after the policy change; the last two columns report the mean difference between these two groups, and the corresponding p-value from the t-test that the means are equal.

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. 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 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

[17] 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.

[18] Appendix Figures A.2a, A.2b, and A.2c 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 initially, apparently as employees overcome initial procrastination.

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.

Table 3: Descriptive Statistics

	Pre: Hired 2005-2009		Post: Hired 2010-2014		Diff.	p-val
	Mean	SD	Mean	SD		
Initial choice, DB or DC with mandatory contributions:						
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 DC plan with mandatory contributions:						
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 > 0)	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
S&P 500 Index, monthly %Δ	0.014	0.033	0.010	0.021	-0.004	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 S&P index.

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 Table 3.[19] 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 it rose in 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. Appendix Table A.1 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 considerably 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 further shows 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.

We first note that there is little change in the proportion of new employees who select the DB plan. 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. One potential concern is that both the benefit cut in the DB plan and the increase in mandatory DC contributions might change the composition of employees selecting the DC plan. For example, perhaps some employees with greater preferences for retirement saving might select the DC plan instead of the DB plan if they are concerned about future state funding of DB benefits. Table 3 shows that the share choosing the DB plan, which is about 20% throughout, fell a little after 2010, by 1.6 pp, while controlling for demographics and job characteristics, we estimate a statistically significant increase of 2.3pp in the share choosing the DB plan. These results suggest at most minor compositional changes and thereby assuage concerns regarding selection into DB plans .

Next, we compare voluntary contributions to the retirement plans (summing together both the 403(b) and 457 options) before and after 2010. These statistics are adjusted for two other changes, described earlier, in retirement plan structure during our time period.

[19]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.

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 participants with very small contribution levels that are consistent with contributing only enough to receive the cash match as not making voluntary contributions; this includes employees making an active choice before 2009. Doing this avoids distorting our participation statistics upwards and our contribution statistics downwards, by excluding employees who may not be making an active choice after the default was instituted. 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.[20] Additionally, in 2013, Roth versions of the voluntary retirement plans become available. Maximum contribution limits, which could then be split between the tax-deferred and Roth accounts, remained the same, and we sum together both afterward, after adjusting the latter for an imputed employee marginal tax rate.[21] 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 a voluntary DC plan at the end of their first nine months of employment, before and after 2010, and Figure 3a shows the same by month. 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.[22] 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

[20] Since we observe not dollars contributed but rather contributions as a percent of salary, we use the upper bound of the salary bin (which is chosen to be conservative) 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 contributing only enough for the cash match and record this contribution as zero. Since the salary bins are relatively narrow and the cash match amounts are small, the estimates are not sensitive to this definition.

[21] 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 noticeably alter contribution amounts.

[22] 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.

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 split the sample by salary, since liquidity constraints may bind more for lower earners. Appendix Table A.1 shows descriptive statistics for the same split, above and below the approximate median salary of $75,000. On average, higher-salaried faculty are more likely to make voluntary contributions, and save more conditional on contributing.

5 Empirical Strategy

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.

5.1 Identification

The ideal experiment to estimate employee responses would be to randomly require some employees to make mandatory contributions while exempting others. Our setting differs from this ideal setup in that all the new faculty hired 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, as in 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. 2022).

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, reducing concerns that changes in hiring practices or in the selection of faculty accepting job offers changed unobservably at the same time.

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

(a) Voluntary participation rate

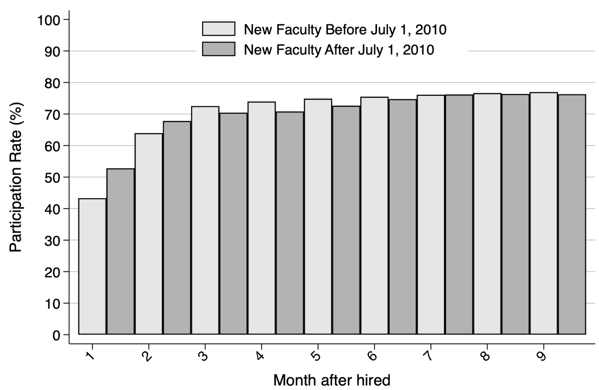

(b) Voluntary contribution rate, if positive

(c) Voluntary contribution rate

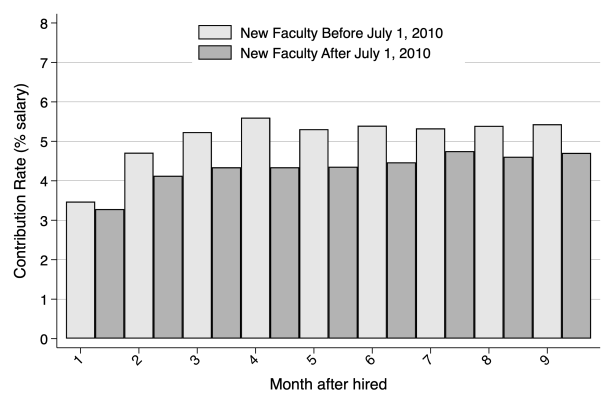

Notes: Figure plots mean participation rates and contribution rates by month since the month faculty are hired. Administrative data on faculty at a large public university, statistics from their first nine months of hire, which covers their first year excluding June through August, when many faculty 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.

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 quite 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. Nor were there substantive 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 extract a sample with similar education and earnings and find 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.2.

As a supplementary strategy, we treat academic staff as a control group when examining voluntary contributions, to difference out any cohort-specific changes in new hires 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.[23] 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.

[23] 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

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 pp 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 during or at the end of the first nine months of employment.

We estimate equations of the following form using OLS:

$$
\begin{aligned}
y_{it} = \alpha &+ \beta_1 Female_i + \beta_2 Single_{it} + \beta_3 Married_{it} + \beta_4 Fulltime_{it} \\
&+ \beta_5 Stock_t + \beta_6 \Delta Stock_t + \sum_a \beta_a \mathbb{1}[Age_i = a] + \sum_j \beta_j \mathbb{1}[Income_i = j] \\
&+ \beta_p Post_i + \sum_k \beta_k \mathbb{1}[Tenure = k] + \sum_t \beta_t \mathbb{1}[Tenure = t] \times Post_i + u_{it} \qquad (4)
\end{aligned}
$$

The key explanatory variables in equation (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 (contributions 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 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 regressions, we control for demographic and employment characteristics that might influence retirement contributions. The demographic characteristics, as discussed in the previous section, are *Female* for gender, *Married* and *Single* for marital status (where the omitted category is unknown status, since this information is reported incompletely), and *Age* for the age ranges that we observe for our sample. The employment characteristics are *Income*, annual salary measured in thousands of 2018 dollars and imputed at 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 both the level and monthly change in the S&P 500 index (*Stock* and Δ*Stock*) measured on the first day of the calendar 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 in the counterfactual, since the full crowd-out benchmark predicts reductions on the intensive margin (but with no welfare loss) for the former and on the extensive margin (with a welfare loss as liquidity constraints tighten) for the latter. 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 liquidity constraints.

To do this, we analyze the distributions of voluntary contributions. We compare the raw distributions for new hires before versus after the policy change, allowing for a straightforward visual contrast against the full crowd-out benchmark. We supplement these comparisons with 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

We now 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. Our benchmark theoretical model, summarized in Table 2 and Figure 2, 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; and 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, primary interest lies with the responses along the full distribution of voluntary contributions, for low versus high earners.

Before presenting our main results, we first establish that there was little change in the share of employees choosing the DC plan. 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 2.6 pp reduction, as shown in Appendix Table A.3, is small relative to baseline DC participation of 79%. This result also contradicts concerns that employees with stronger preferences for saving would be more likely to select the DC plan after the cut to the DB plan. Therefore, we have little ground for concern about compositional changes among DC plan participants. For robustness, we also estimate voluntary contributions using all faculty, including those who chose the DB plan, and find similar results as our main estimates (Appendix Figure A.6 and Appendix Table A.6).

6.1 Full-Sample Estimates

We now present our estimation results in Table 4 for voluntary participation and contribution rates for the full sample. The table reports 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 statistically insignificant starting in the 3rd month and often slightly positive, though theory predicts a reduction in participation.[24] Among those who participate, the contribution rate declines significantly by about 2.5 pp, and the lower bound of the 95 percent confidence intervals show an average decrease of around 4 pp.[25] Even this lower bound estimate is far from the potential for crowd-out on the participation margin, since about 50 percent of prior participants contributed less than the new required threshold in the pre-treatment period.

Estimates for the voluntary contribution rate by the full sample give a complete picture of the crowd-out response. These appear in the final column of Table 4 for the 3rd, 6th, and 9th months after hire, and in Figure 4c for each month. For the full sample, the average reduction is statistically significant and generally between 1.5–2 pp.

This evidence of limited crowd-out is supported by an event-study analysis that appears in Appendix Figure A.8 (which also shows university staff, whom we use later as a control group). While the estimates are somewhat noisy, likely reflecting our small sample size, they show a reduction in average voluntary contributions that is considerably less than 3.5–5 pp. 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

[24]In the first two months, participation is substantially higher in the *Post* period, while the contribution rate conditional on participating is quite similar across months. As we showed earlier in the raw statistics and later highlight in the distributional analysis, the unconditional mean participation rate is in fact lower post-2010 for new employees nine months after hire.

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

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	0.011	(0.025)	-2.809	(0.735)	-1.743	(0.536)
Post-2010 x Tenure month 6	0.013	(0.026)	-2.771	(0.738)	-1.699	(0.552)
Post-2010 x Tenure month 9	0.011	(0.026)	-2.391	(0.767)	-1.468	(0.585)
Age, relative to $[25, 30)$						
$[30, 35)$	-0.033	(0.033)	0.464	(0.744)	0.003	(0.522)
$[35, 40)$	-0.043	(0.034)	1.111	(0.811)	0.425	(0.578)
$[40, 45)$	-0.082	(0.036)	1.507	(0.916)	0.499	(0.654)
$[45, 50)$	-0.133	(0.042)	2.844	(1.173)	1.033	(0.822)
$[50, 55)$	-0.115	(0.046)	6.320	(1.729)	3.552	(1.247)
$[55, 60)$	-0.087	(0.049)	2.349	(1.487)	0.789	(1.076)
$[60+)$	-0.227	(0.0578)	1.416	(1.795)	-0.739	(1.164)
Salary, relative to $[20k - 25k)$						
$[25k - 29k)$	0.171	(0.0610)	0.820	(1.658)	1.648	(0.856)
$[30k - 35k)$	0.094	(0.0635)	0.515	(1.517)	1.164	(0.722)
$[35k - 39k)$	0.173	(0.0629)	-1.263	(1.391)	0.367	(0.686)
$[40k - 45k)$	0.193	(0.0642)	1.782	(1.618)	2.203	(0.866)
$[45k - 49k)$	0.278	(0.0585)	0.723	(1.523)	2.158	(0.914)
$[50k - 55k)$	0.283	(0.0582)	0.640	(1.505)	2.305	(0.906)
$[55k - 59k)$	0.323	(0.0624)	1.778	(1.704)	3.234	(1.069)
$[60k - 65k)$	0.381	(0.0577)	1.489	(1.656)	3.315	(1.062)
$[65k - 69k)$	0.396	(0.0588)	2.058	(1.563)	3.790	(1.006)
$[70k - 75k)$	0.390	(0.0602)	-0.389	(1.358)	2.051	(0.782)
$[75k - 79k)$	0.383	(0.0664)	0.973	(1.908)	3.056	(1.262)
$[80k - 85k)$	0.349	(0.0695)	2.888	(2.011)	4.091	(1.406)
$[85k - 89k)$	0.470	(0.0601)	1.231	(1.628)	3.687	(1.113)
$[90k - 95k)$	0.393	(0.0628)	2.631	(1.733)	4.188	(1.102)
$[95k - 99k)$	0.418	(0.0624)	2.105	(1.763)	3.657	(1.169)
$[100k - 120k)$	0.472	(0.0495)	5.205	(1.361)	6.851	(0.773)
$[120k - 140k)$	0.438	(0.0602)	4.246	(1.813)	5.884	(1.293)
$[140k+)$	0.441	(0.0550)	-0.358	(1.369)	3.512	(0.767)
N	17,961		12,725		17,961	

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.

longer-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 are smaller.

The results therefore indicate crowd-out rates far below what is predicted by standard economic theory. Full crowd-out implies a 3.5–5 pp reduction, yet the lower bounds of the 95 percent confidence intervals reject even a 3.5 pp reduction in contribution rates, while the participation response is far too small to explain less-than-full crowd-out. 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% for mandated employee contributions.[26]

6.2 Distributional Analysis by Earnings Level

Next, we use quantile plots to evaluate crowd-out effects. We display the raw distribution of contribution rates of employees 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 of 5 pp by subtracting this amount (or the 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 differences between groups at lower contribution rates.[27]

These quantile plots show a modest participation response among low-contributing low earners, and a modest full 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 rises from about 30% before the policy to just over 40% after. Yet, this is far below the drop in participation expected under the full crowd-out benchmark, in which over 80% of lower-earning employees who contributed prior to the policy change 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.

[26] Dividing the 1.468 pp reduction in voluntary contributions after 9 months from Table 4 by 5pp yields a 29.4% crowd-out rate.

[27] We also examine the probability of making the maximum 403(b) contribution (with too few individuals participating in the 457 to make this a meaningful part of the analysis), and find a small reduction of 3 pp (Appendix Figure A.9), which is about 30 percent of the baseline mean.

Figure 4: Event study regression estimates

(a) Voluntary participation rate

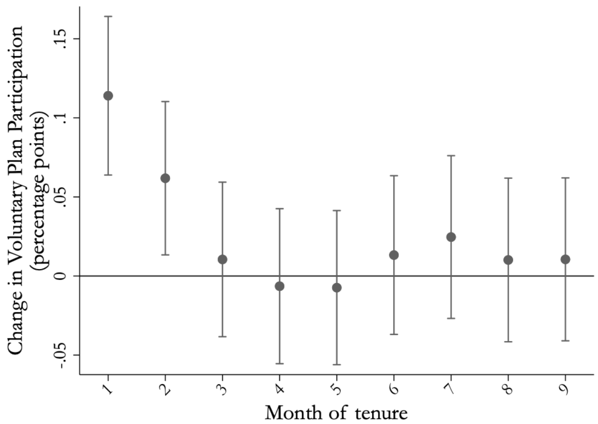

(b) Voluntary contribution rate, if positive

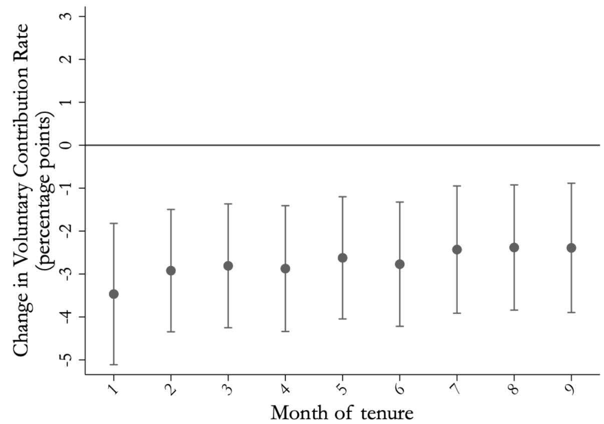

(c) Voluntary contribution rate

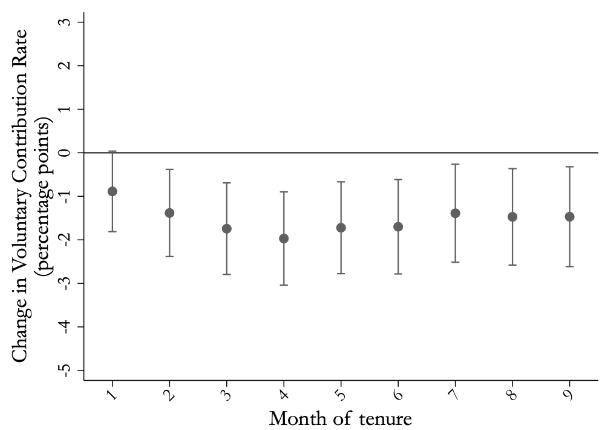

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the event study regressions corresponding to equation 4. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. Additional 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.

The patterns differ among employees with salaries over $75,000, in the right panel of Figure 5. 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.[28]

To formally examine these distributional differences, we estimate 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. Appendix Figure A.11 presents the coefficient estimates on $Post$ from each percentile, with the shaded region representing 95 percent confidence intervals.

The quantile regressions confirm the patterns documented in Figure 5. For employees below $75,000 in salary, the changes are small, negative, and precisely estimated for 80 percent of the distribution of contributions, and then become noisier at the higher end (Figure A.11a). By contrast, for employees above $75,000 in salary, reductions of around 5 pp occur between the 55th and 85th quantiles, as well as at some higher quantiles (Figure A.11b).

In summary, these results provide evidence of incomplete crowd-out for employees, with small shares responding with full crowd-out and a large share not responding noticeably. The overall shares of employees with each type of response is 17% with a full 5 pp reduction in saving (or active high savers, as we termed them in the Introduction); 8% ceasing to participate (active low savers); and 50% displaying little or no crowd-out (passive savers).[29] The percentage of active high savers in our sample is very similar to that found in the analysis of a considerably smaller mandatory contribution in Denmark in Chetty et al. (2014). Moreover, the active high savers are largely higher earners (with little evidence of crowd-out at the low end of the distribution of contribution rates), and the active low savers are largely lower earners (with a larger share of low earners exhibiting passive responses).[30]

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

[29] The remaining 25% of the sample did not make voluntary contributions in the first place.

[30] Appendix Figure A.10 shows similar patterns when splitting the sample by age, which is highly correlated with salary.

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

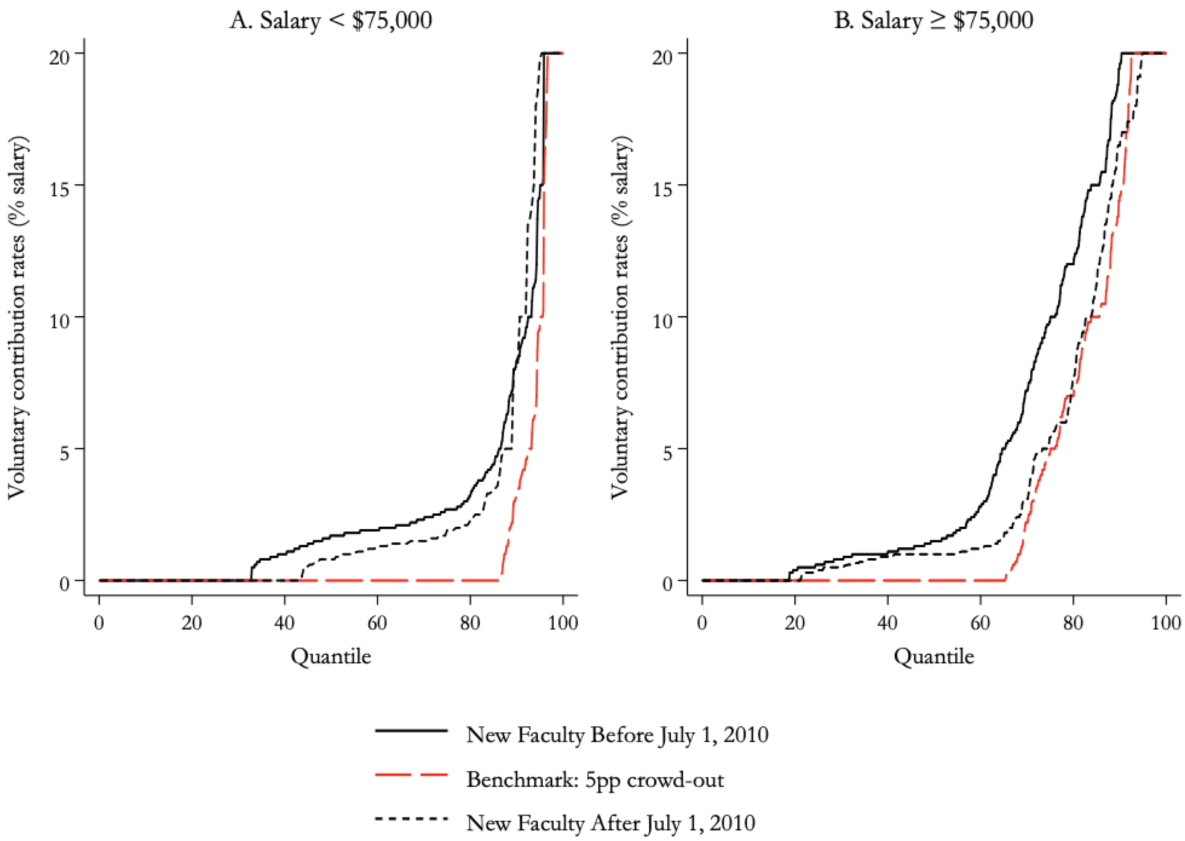

Notes: Figure plots the ordered values of voluntary contribution rates in the 9th month of tenure against quantiles, separately for faculty hired before July 1, 2010 and faculty and faculty hired after this date. Panel (a) includes faculty with annual salaries less than $75,000 and Panel (b) includes 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.

6.3 Robustness

We implement several robustness checks and find similar results: we continue to strongly reject full crowd-out. First, we estimate equation (4) without any of the control variables for demographics, job characteristics, and financial market returns. Appendix Figure A.5 shows that the response is slightly more modest, with most estimates at 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, as employees hired then may have been leery of voluntary contributions. Yet, we continue to reject a 5 pp reduction based on the lower bound of the 95% confidence intervals (Appendix Figure A.3). We also find incomplete crowd-out if we drop the first two years and the last two years of the sample, restrict attention to between 2007 and 2012 (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.2 show similar patterns for the first 27 months as the first 9 months.

We then 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.7 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 not directly affected by the policy change, since they do not participate in the mandatory DC plan, to conduct a difference-in-difference analysis.[31] Using a control group helps to address concerns about changes across cohorts in preferences, composition, or information about retirement saving, whether driven by changes at the employer or in the economy. While saving rates between staff and faculty differ in levels, the identifying assumption behind the difference-in-difference analysis relies on the assumption of parallel trends and does not require equal levels of saving before 2010.

We find essentially no change in the staff voluntary contribution rate, either in 2010

[31] 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 gradually over time. Appendix Figure A.8 plots average contribution rates separately for faculty and staff, and split by whether the employee was hired before or after the policy change. The rates 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.[32] 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 for affected faculty. The lower bound of the 95% confidence interval is a -1.6pp reduction, which is approximately the 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 support for our finding that crowd-out is limited, on average, and is well below the neoclassical benchmark.

6.4 Consumption Equivalent Calculations

To illustrate the potential implications of incomplete crowd-out, we parameterize the two-period model from Section 3, which predicts full crowd-out except for those whose liquidity constraints become more binding. We quantify the "consumption equivalent", which is the dollar amount that a person would be willing to forgo each period to avoid making mandatory contributions—or alternately, that they would be willing to pay to be required to make them. While we know the magnitude of active versus passive responses, we do not know whether someone saving either passively or actively is being nudged towards or away from their optimal savings rate. Therefore, we consider some extreme cases and indicate gaps in our knowledge that future research can explore so as to move towards a welfare analysis of mandatory savings policies.

Appendix A provides details of the equations used to compute the consumption equivalents. A critical assumption is the extent to which people were saving optimally before the policy. We introduce a parameter that governs the extent to which people were already optimizing without mandatory contributions, and calculate the consumption equivalents under a range of values for this parameter to illustrate how magnitudes vary with different assumptions. At one extreme, if all employees were already saving optimally, then both partial crowd-out by active savers who stop participating (but cannot reduce saving by more) and also lack of crowd-out by passive savers (who should not save more) brings the average loss across the full sample to about $1,765 per person in each period

[32] 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.

(equal to a 3.0% loss in consumption). At the other extreme, if we assume that everyone was initially undersaving and the policy brings them to their optimum, then the average benefit is about $2,225 (equal to a 4.0% gain in consumption).

In reality, some employees were likely under-saving without the policy while others were already saving optimally. Some may even have been over-saving, in which case the costs would be larger for those who do not offset the higher forced contributions with lower voluntary contributions, while the policy may induce some who were undersaving to oversave, offsetting some or all of their gains. The average effects of the policy will be a combination of the relative shares of each group along with the consumption equivalents shown in Figure A.12. As an example, if those who exhibit full crowd-out were saving optimally and those who exhibit no crowd-out or partial crowd-out are now saving optimally, the gains to the latter exceed the costs to the former by about $550, on average.

There are two key qualitative takeaways from this exercise. First, the benefits or costs of the policy are not negligible, which is not surprising given the large magnitude of the increase in mandatory contributions. Second, we find that the consumption equivalents are larger in magnitude, in one direction or the other, for lower earners because they exhibit less crowd-out. Mandatory contributions are therefore likely to have a larger impact (positive or negative) on employees with lower salaries, and future research should explore how the need for liquidity balances against the need for commitment in this group.

7 Discussion

Even though the goal of requiring employees to participate in retirement savings plans, including public or private pension plans, is presumably to induce more saving, standard economic theory predicts crowd-out: higher required contributions should lead employees to reduce their voluntary contributions. 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 less, to new employees hired before, we detect average reductions of voluntary contribution rates of around 1.5 pp. Our confidence intervals rule out the 3.5-5 pp reductions predicted by theory.

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. Otherwise, for half of the sample, we observe little or no crowd-out in voluntary contributions. On average, our results suggest that mandatory retirement contributions are likely to raise total saving rates.

It is important to acknowledge that our results pertain to a single setting. While employees in our sample have relatively high incomes and education levels, this apparently does not forestall a great deal of passive behavior. A limitation of our analysis, like many in this area, is that it only records savings in workplace plans, which prevents us from studying adjustments along other margins.

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 (Choukhmane 2021) 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.

The welfare effects of the policy depend on the extent to which people were optimally saving without the policy. Whether or not Americans are saving adequately for retirement is a long-debated question, however, and our study does not yield an answer to this question. Instead, our calculations illustrate that the costs and benefits can be sizable in magnitude, and are larger, whether they are gains or losses, for the group of lower-income employees because they exhibit less crowd-out than higher-income employees. 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

Aguila, Emma, "Personal Retirement Accounts and Saving," *American Economic Journal: Applied Economics*, 2011, *3* (4), 1–24.

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

Attanasio, Orazio and Agar Brugiavini, "Social Security and Households' Saving," *Quarterly Journal of Economics*, 2003, *118* (3), 1075–1119.

— **and Susann Rohwedder**, "Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom," *American Economic Review*, 2003, *106* (4), 1499–1521.

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

— **,** — **, David Laibson, Brigitte Madrian, and William Skimmyhorn**, "Borrowing to Save? The Impact of Automatic Enrollment on Debt," *Journal of Finance*, 2022, *77*, 403–447.

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," 2021.

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

Engelhardt, Gary and Anil Kumar, "Pensions and Household Wealth Accumulation," *Journal of Human Resources*, 2011, *46* (1), 203–226.

Fadlon, Itzik and David Laibson, "Paternalism and Pseudo-Rationality," 2019.

Feldstein, Martin, "Social Security, Induced Retirement, and Aggregate Capital Accumulation," *Journal of Political Economy*, 1974, *82* (5), 905–926.

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

Hubbard, Glenn and Kenneth L. Judd, "Liquidity Constraints, Fiscal Policy, and Consumption," *Brookings Papers on Economic Activity*, 1986, *17* (1), 1–60.

Hubbard, R. Glenn, "The Effects of Pensions on Household Wealth: A Reevaluation of Theory and Evidence," *Journal of Political Economy*, 1998, *106* (4), 706–723.

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

Lachowska, Marta and Michal Myck, "The Effect of Public Pension Wealth on Saving and Expenditure," *American Economic Journal: Applied Economics*, 2018, *10* (3), 284–308.

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

Leive, Adam, Leora Friedberg, and Brent Davis, "Overpaying and Undersaving? Correlated Mistakes in Retirement Saving and Health Insurance Choices," 2022.

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.

Michaud, Pierre-Carl Hurd Michael and Susann Rohwedder, "The Displacement Effect of Public Pensions on the Accumulation of Financial Assets," *Fiscal Studies*, 2012, *33* (1), 107–28.

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

Munnell, Alicia, "Private Pensions and Savings: New Evidence," *Journal of Political Economy*, 1976, *84* (5), 1013–1032.

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

Renata, Jappelli Tullio Bottazzi and Mario Padula, ""Retirement expectations, pension reforms, and their impact on private wealth accumulation," *Journal of Public Economics*, 2006, *90* (12), 2187–2212.

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.

Yang, Tzu-Ting, "The Effect of Workplace Pensions on Household Saving: Evidence from a Natural Experiment in Taiwan," *Journal of Risk and Insurance*, 2020, *87* (1), 173–194.

Internet Appendices [Not for Publication]

Internet Appendix A: Supplementary Analyses [Not for Publication]

Additional Details of Demographics: As additional evidence on balance in demographics across cohorts, the figure below plots means for faculty by year of hire. The vertical line at 2010 denotes the year of the policy change. Consistent with the balance results shown in Table 3, there are not visible breaks in demographics or salary among new hires before and after the policy change.

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

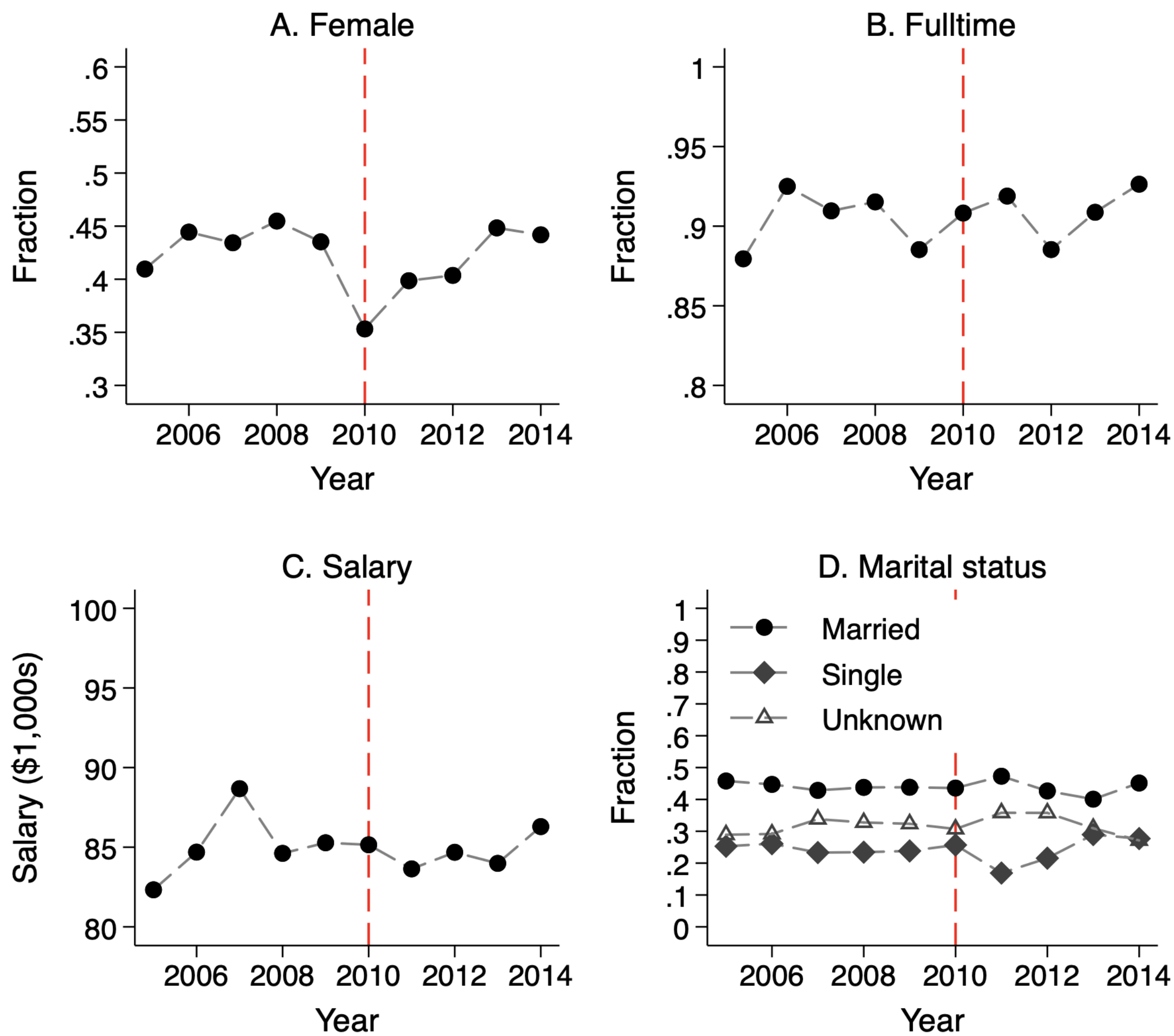

Notes: Administrative data on faculty at a large public university, statistics from their first month of hire. Figures plot the averages among employees hired each year from 2005 to 2014.

Table A.1 presents descriptive statistics separately for employees with annual salaries below or above $75,000. For both of these groups, demographics are balanced between new faculty hired before the policy change vs. afterwards.

Table A.1: 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.

Analysis using the Survey of Consumer Finances: To assess whether there are any national-level trends in saving during our sample period, we calculate mean participation and contribution rates in voluntary plans using the Survey of Consumer Finances (SCF). Table A.2 below reports means weighted using survey sample weights from the 2004, 2007, 2010, 2013, and 2016 SCF surveys, which are conducted every three years. 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). Mean rates are roughly flat during this period among this sample of employees. This pattern suggests there were no aggregate-level changes in saving rates that coincided with the policy change at the university we study.

Table A.2: Mean Participation and Contribution Rates in the SCF

Survey Year	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%

Choice of Mandatory DB vs. DC plan: As described in the text, employees face a one-time, irrevocable choice between a DB plan and the DC plan with mandatory contributions. The default is the DB plan: if they do not actively choose the DC plan within 60 days of starting employment, they are enrolled in the DB plan. Table A.3 shows there is no evidence of a large change in the probability of choosing the DB plan after 2010. We estimate a 2.3pp increase in the probability of DB that is not statistically significant (SE = 2.1pp). Since the choice is irrevocable, we include one observation per employee, corresponding to their 9th month of hire.

Table A.3: OLS Regression Estimates, DB Plan Choice

	Estimate	SE
Post-2010 policy change	0.026	(0.021)
Age, relative to $[25, 30)$		
$[30, 35)$	0.025	(0.030)
$[35, 40)$	0.047	(0.031)
$[40, 45)$	0.054	(0.034)
$[45, 50)$	0.051	(0.039)
$[50, 55)$	0.044	(0.042)
$[55, 60)$	0.074	(0.048)
$[60+)$	0.113	(0.050)
Income, relative to $[20k - 25k)$		
$[25k - 29k)$	-0.080	(0.045)
$[30k - 35k)$	-0.008	(0.049)
$[35k - 39k)$	0.081	(0.049)
$[40k - 45k)$	0.267	(0.052)
$[45k - 49k)$	0.220	(0.050)
$[50k - 55k)$	0.121	(0.048)
$[55k - 59k)$	0.020	(0.054)
$[60k - 65k)$	0.001	(0.051)
$[65k - 69k)$	0.012	(0.053)
$[70k - 75k)$	-0.027	(0.053)
$[75k - 79k)$	-0.026	(0.058)
$[80k - 85k)$	-0.112	(0.053)
$[85k - 89k)$	-0.109	(0.053)
$[90k - 95k)$	-0.135	(0.050)
$[95k - 99k)$	-0.130	(0.051)
$[100k - 120k)$	-0.190	(0.038)
$[120k - 140k)$	-0.073	(0.052)
$[140k+)$	-0.135	(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. Regressions also include indicators for sex, fulltime employee, marital status, the level and percentage change of the S&P 500 on the 1st of each month, and a constant.

Crowd-out patterns over the first 3 years: To assess whether there may be more crowd-out over a longer horizon, Figure A.2 plots regression results extending to 27 months (corresponding to 3 academic years, excluding summer months). We continue to reject a 5pp reduction throughout the first three years, and also reject a 3.5pp reduction except in the third year. The point estimates reveal a 2pp decline. To assess the possibility of differential exits over time, Table A.4 presents descriptive statistics and balance tests on the composition of employees who remain employed after 18 and 24 months. At 18 months, the difference on the omnibus test is marginal at the 5% level, driven by changes in marital status variables. We fail to reject the null of no change in composition after 24 months, though the sample is also smaller. Taken together, there is not strong evidence that the composition of those who remain employed is systematically different for employees hired before 2010 versus afterwards.

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 shows 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.

Figure A.2: Voluntary participation and contribution rates by tenure month, newly hired faculty, first 3 years

(a) Voluntary participation rate

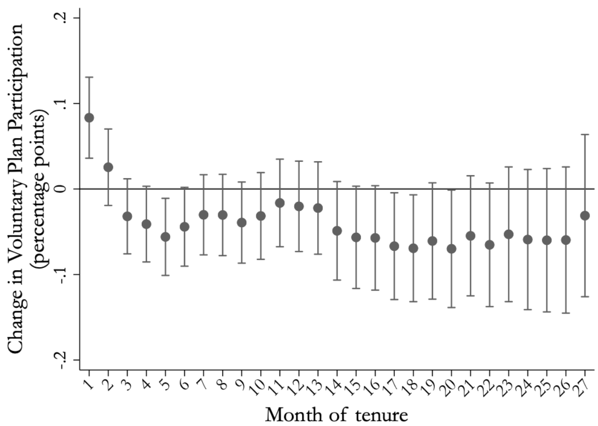

(b) Voluntary contribution rate, if positive

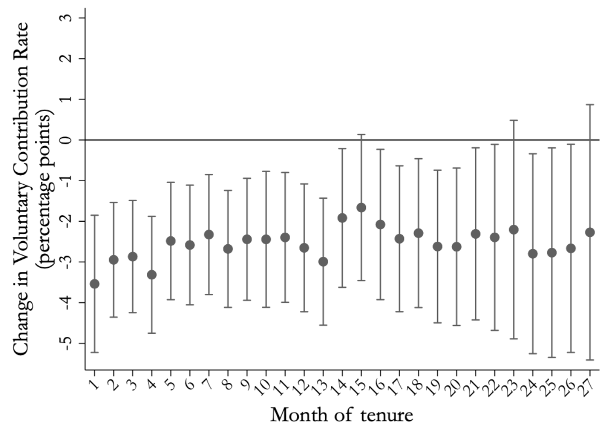

(c) Voluntary contribution rate

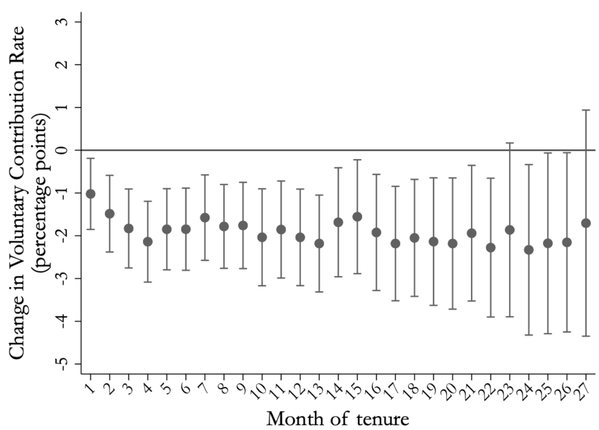

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the regressions corresponding to equation (4). Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. The 27th month of tenure corresponds to the end of the 3rd academic year since being hired. 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.

Robustness to Sample and Regression Specification: We conduct several robustness checks to help validate our main results. In each case, we present the figures corresponding to Figure 4a, Figure 4b, and Figure 4c which compare rates for the first nine months for new hires starting after the policy change versus new hires who started before the change. This period covers their first year excluding June through August, when many faculty do not receive a salary.

1. Exclusion of Great Recession period: Figure A.3 shows that patterns are similar when excluding employees hired between September 2008 and December 2009, inclusive. We continue to find participation estimates near zero and contribute rate changes of similar magnitudes as the main results.

2. Exclusion of end of sample period: Figure A.4 shows that we continue to find incomplete crowd-out when excluding employees hired prior to 2007 or after 2012 to restrict attention to the period surrounding the policy change.

3. Regressions Excluding Controls: Figure A.5 displays estimates that exclude controls. Contribution rates are now estimated to decline by about 1 percentage point, and we statistically reject declines of 2 percentage points. These patterns indicate even less crowd-out compared to the main estimates.

4. Including faculty who choose mandatory DB plan: The results are robust to including all new faculty, including those who select the DB plan. This is important because although the 5% contribution increase only applied to the DC plan, there may be a concern that focusing only on employees who choose the DC selects on an outcome. This sample restriction might cause problems due to changes in sample composition. While we have shown shown that the probability of choosing the DB plan does not change after the policy, we view it as important to also run models that include all new faculty, including those who do not choose the mandatory DC plan. Figure A.6 shows results are very similar to our main results.

Figure A.3: Robustness: Voluntary participation and contribution rates by tenure month, newly hired faculty, excluding Sept 2008 – Dec. 2009

(a) Voluntary participation rate

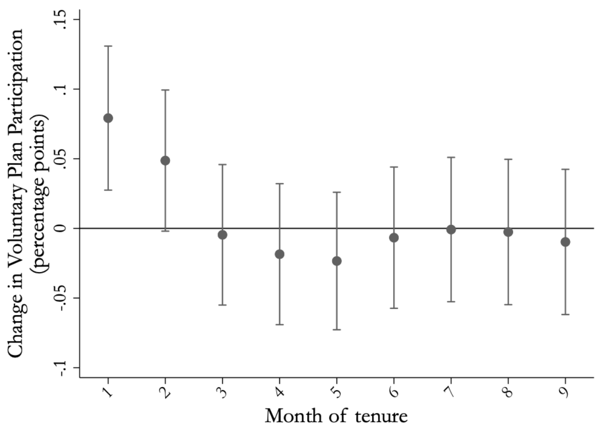

(b) Voluntary contribution rate, if positive

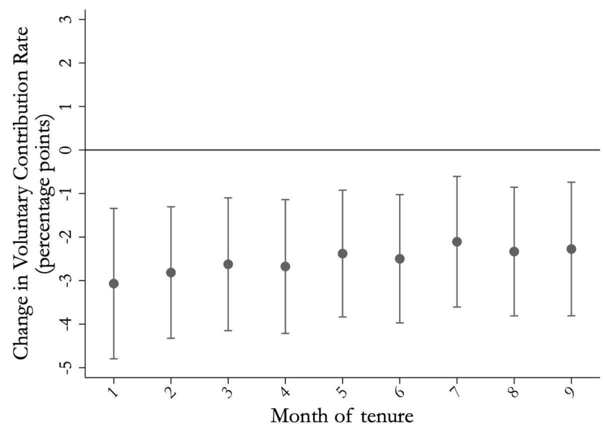

(c) Voluntary contribution rate

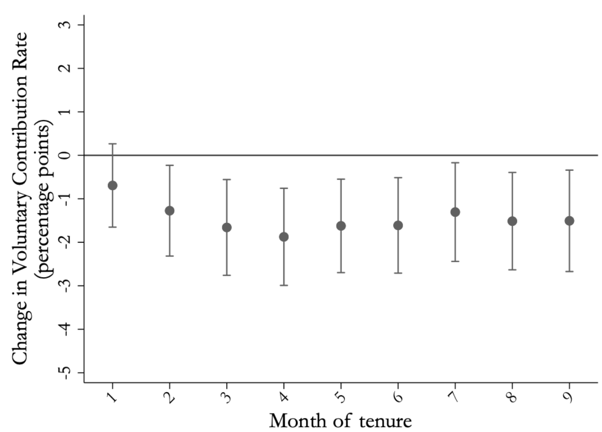

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the regressions corresponding to equation (4). Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. The sample excludes observations from September 2008 to December 2009, inclusive. Regressions include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands as controls.

Figure A.4: Robustness: Voluntary participation and contribution rates by tenure month, newly hired faculty, 2007–2012

(a) Voluntary participation rate

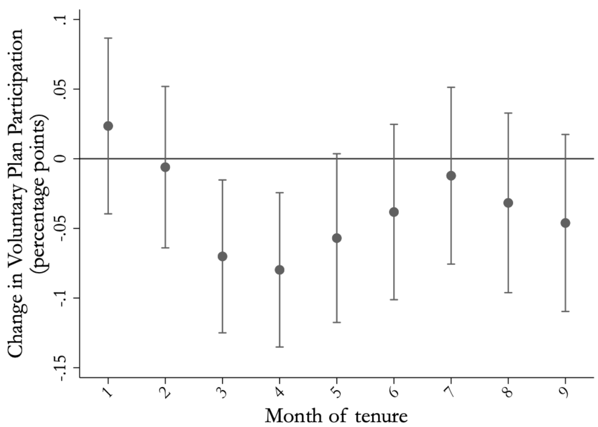

(b) Voluntary contribution rate, if positive

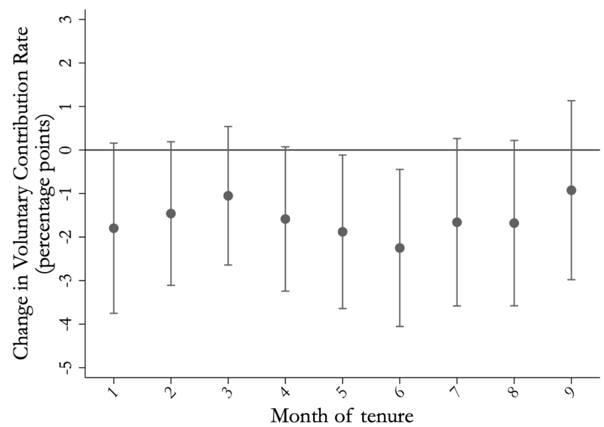

(c) Voluntary contribution rate

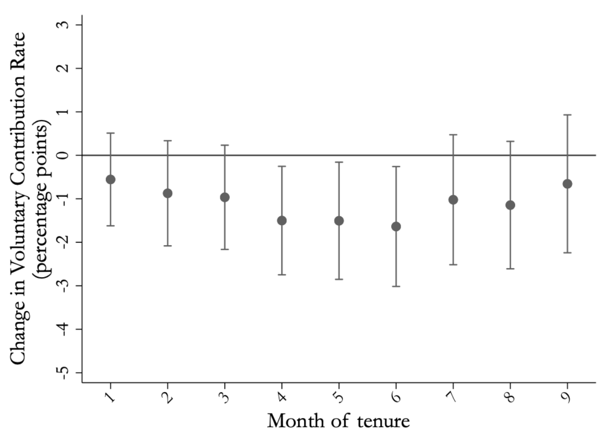

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the regressions corresponding to equation (4). Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. The sample excludes observations from before 2007 and after 2012. Regressions include the monthly S&P 500 index, and indicators for tenure month, female, full-time employee, married, single, and income and age bands as controls.

Figure A.5: Robustness: Voluntary participation and contribution rates by tenure month without controls, newly hired faculty

(a) Voluntary participation rate

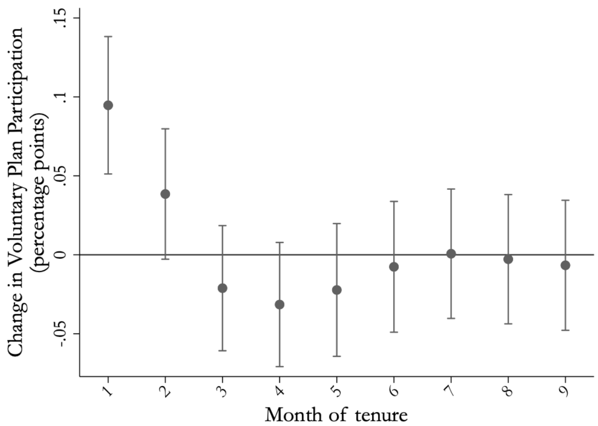

(b) Voluntary contribution rate, if positive

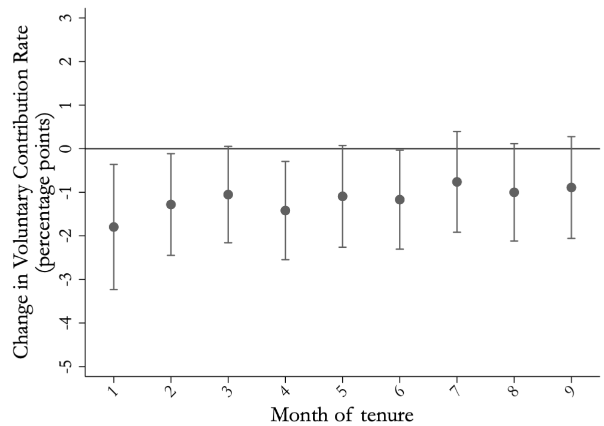

(c) Voluntary contribution rate

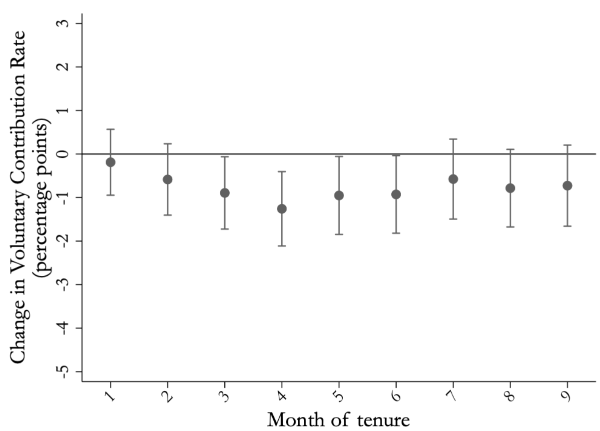

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. Additional controls include indicators for tenure month and a constant but exclude other controls.

Figure A.6: Robustness: Voluntary participation and contribution rates by tenure month, all newly hired faculty

(a) Voluntary participation rate

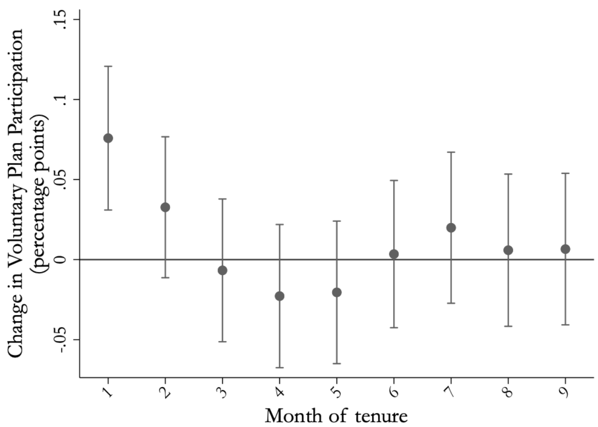

(b) Voluntary contribution rate, if positive

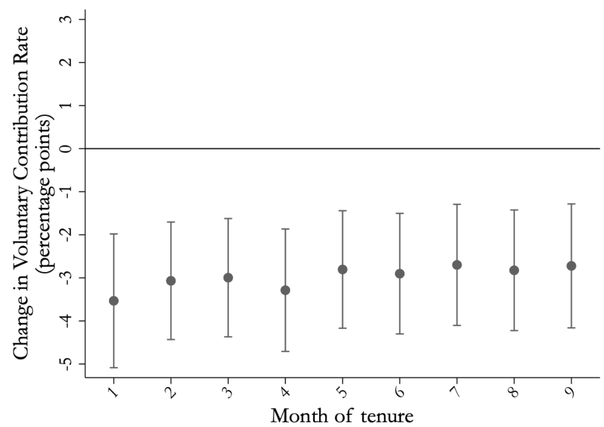

(c) Voluntary contribution rate

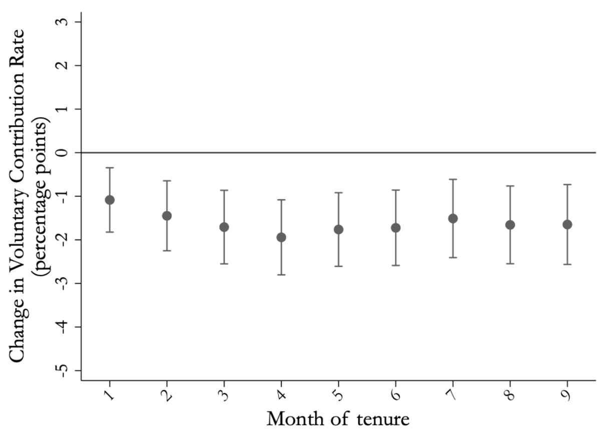

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the regressions corresponding to equation (4). Whiskers denote 95 percent confidence intervals with standard errors clustered by employee. The sample includes 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.

Placebo test: Choice of vendor: Employees have a choice of two vendors (which we label "A" and "B" to protect confidentiality) for their 403(b). Both vendors offer a range of different funds, simple online interfaces, research tools, and other standard features. Fees are similar and low in both. Compared to vendor A, vendor B focuses to a greater extent on active wealth management and the two vendors differ slightly in their investment options. The choice of which vendor should not be impacted by the policy change, but if preferences for saving changed across cohorts, one might expect a change in the choice of vendor. Figure A.7 shows cohorts before and after the change did not differ in the preference for one vendor over the other.

Figure A.7: Placebo Test: Choice of Vendor in Retirement Accounts

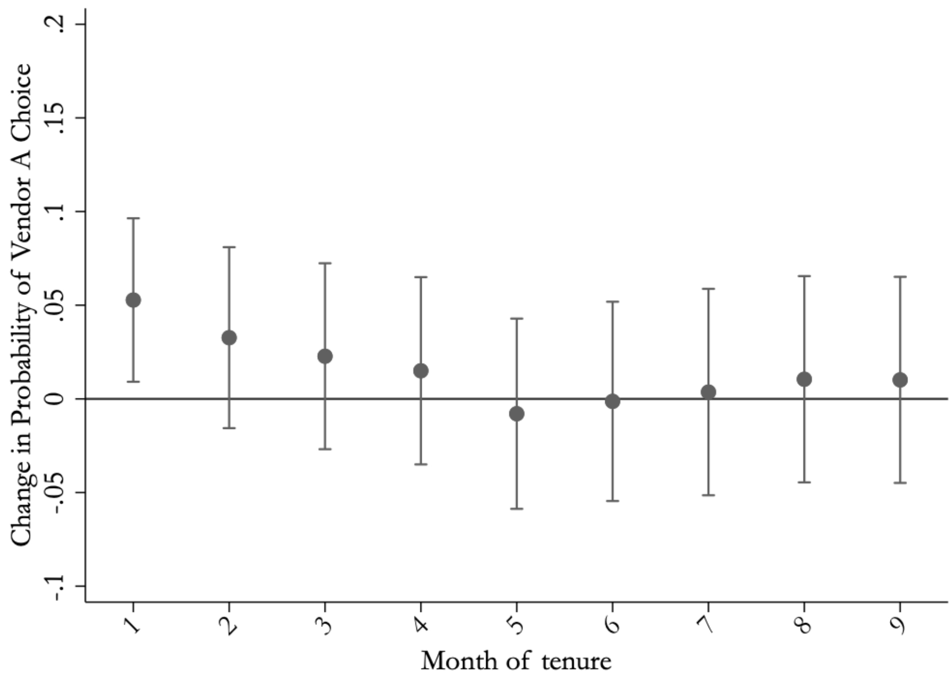

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the regressions corresponding to (4) in the main text. The dependent variable is an indicator for choosing Vendor A (instead of Vendor B) as the account for voluntary 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.

Difference-in-Difference Analysis Using Staff as a Control Group: As a second identification strategy, we use staff as a control group. Staff were not eligible for the mandatory DC plan and so were not affected by the 5% mandatory contribution. They were still able to make voluntary contributions to both 403(b) and 457 accounts, just like faculty. Table A.5 shows descriptive statistics for staff. The average age remained flat at 36 years, while the mean (real) salary increased slightly, from $43,900 to $45,700. Participation rates increased from 46% to 54% but the contribute rate remained similar at 1.8% of salary. Figure A.8 shows mean contribution rates for staff and faculty by year of hire. 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.

We then estimate the following difference-in-difference specification to provide a single measure of crowd-out:

$$y_{it} = \alpha + \gamma_1 Faculty_i + \gamma_2 Post_i + \gamma_3 Faculty_i \times Post_i + e_{it} \tag{A.5}$$

where y_{it} denotes the voluntary contribution rate for employee i in tenure month t. $Faculty_i$ is an indicator for faculty relative to staff, $Post_i$ is an indicator for employees hired after the policy change, and γ_3 represents the coefficient of interest. We also estimate models with the same controls as the primary specification in the main text. Table A.6 shows the regression results. Columns 1 and 2 include faculty who choose the mandatory DC plan, as in our main specification. Columns 3 and 4 include all faculty, including those who choose the DB plan. Columns 1 and 3 exclude controls and columns 2 and 4 include controls. The DD estimate ranges from -0.701 (column 1) to -1.067 (column 4), with each being statistically significant at conventional levels. Estimated crowd-out rates are therefore closer to zero when including staff, compared to the main results.

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 ($1,000s)	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

Figure A.8: Regression-Adjusted Total Contribution Rates by Faculty/Staff and Year of Hire

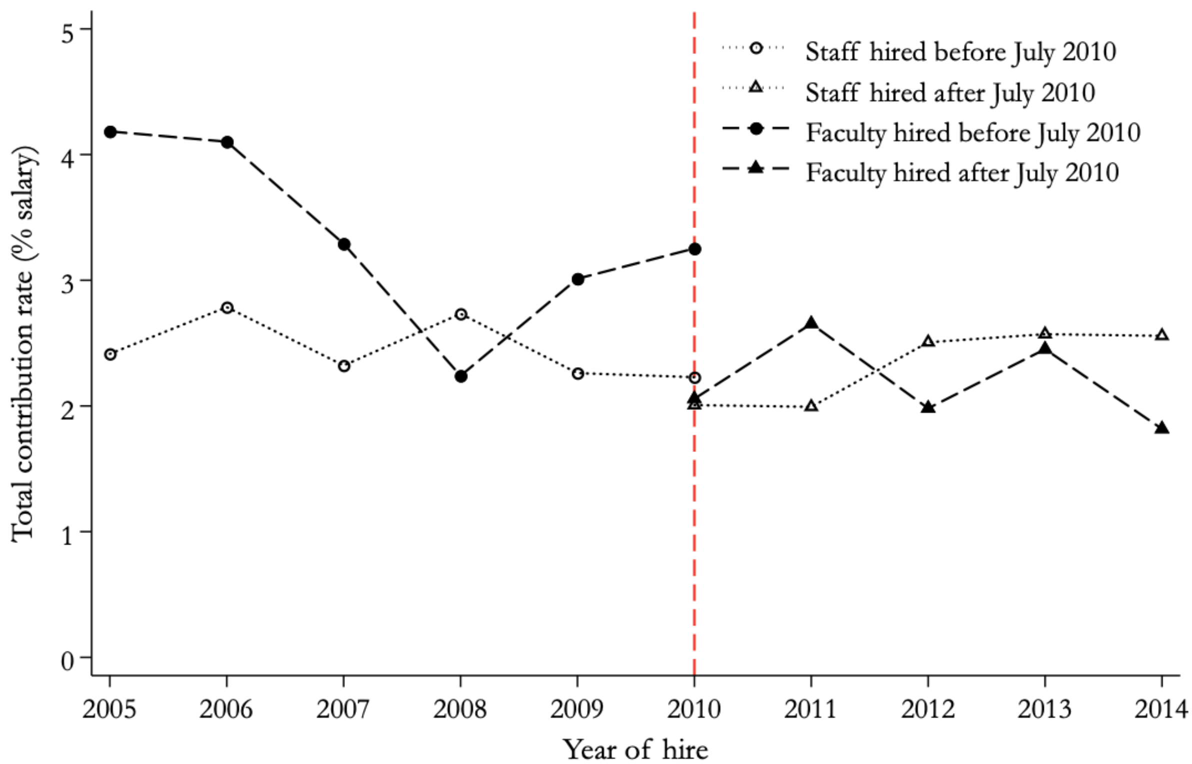

Notes: Figure plots mean total voluntary contribution rates for four groups: faculty hired before the policy change (shaded circles), faculty hired after the policy change (shaded triangles), staff hired before the policy change (hollow circles), and staff hired after the policy change (hollow triangles).

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.

Analysis of maximum contributions: We observe a moderate reduction in the probability of maxing out voluntary contributions in response to the policy change. Figure A.9 shows there is a roughly 3 pp decline in making the maximum voluntary contribution. This result is consistent with crowd-out among employees with high salaries who make high voluntary contributions.

Figure A.9: Probability Contribute at Maximum

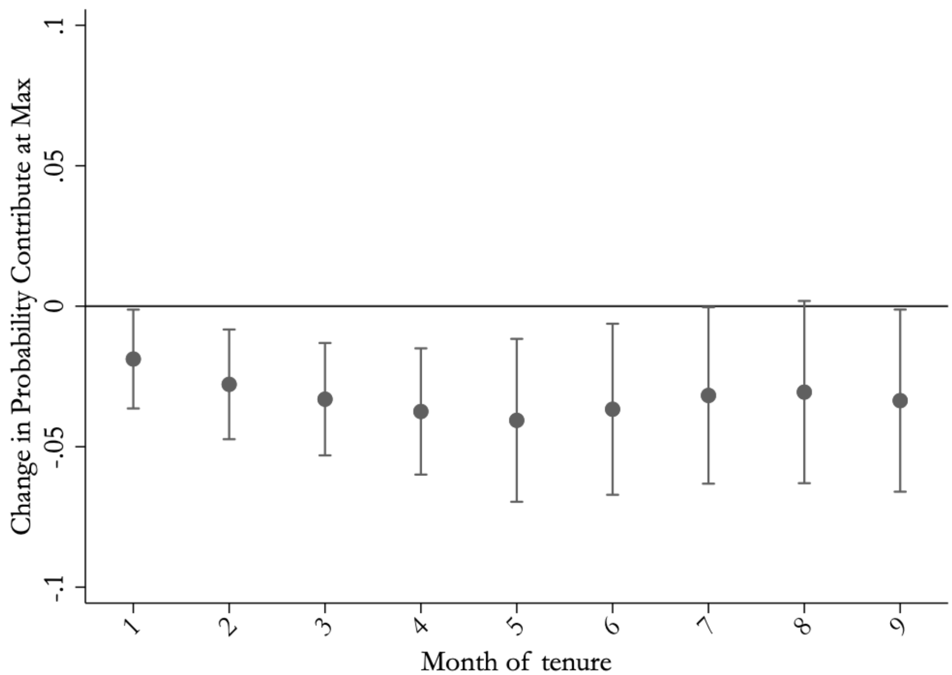

Notes: Figure plots OLS estimates of the interaction between indicators for tenure month and an indicator for being hired after July 2010 from the regressions corresponding to (Equation 4) in the main text. The dependent variable is an indicator for making the maximum contribution to the voluntary plan. Whiskers denote 95 percent confidence intervals with standard errors clustered by employee.

Additional Details of Quantile Analysis: Figure A.10 reproduces the quantile plots in Figure 5 split by age, rather than salary. In several ways, the patterns are similar to the split by income, which is consistent with the high correlation between age and income. In particular, the patterns are similar between those under age 35 and those earning less than \$75,000: there is evidence of reduced participation but otherwise little crowd-out. For those over age 35, there is no sign of a participation response, just as was observed for those earning over \$75,000. There is, however, less evidence of crowd-out and instead mostly passive saving at the higher end of the distribution of contributions for those under 35, unlike the case for higher-salaried employees.

Figure A.10: Quantile plot of voluntary contribution rates, by age and date of hire

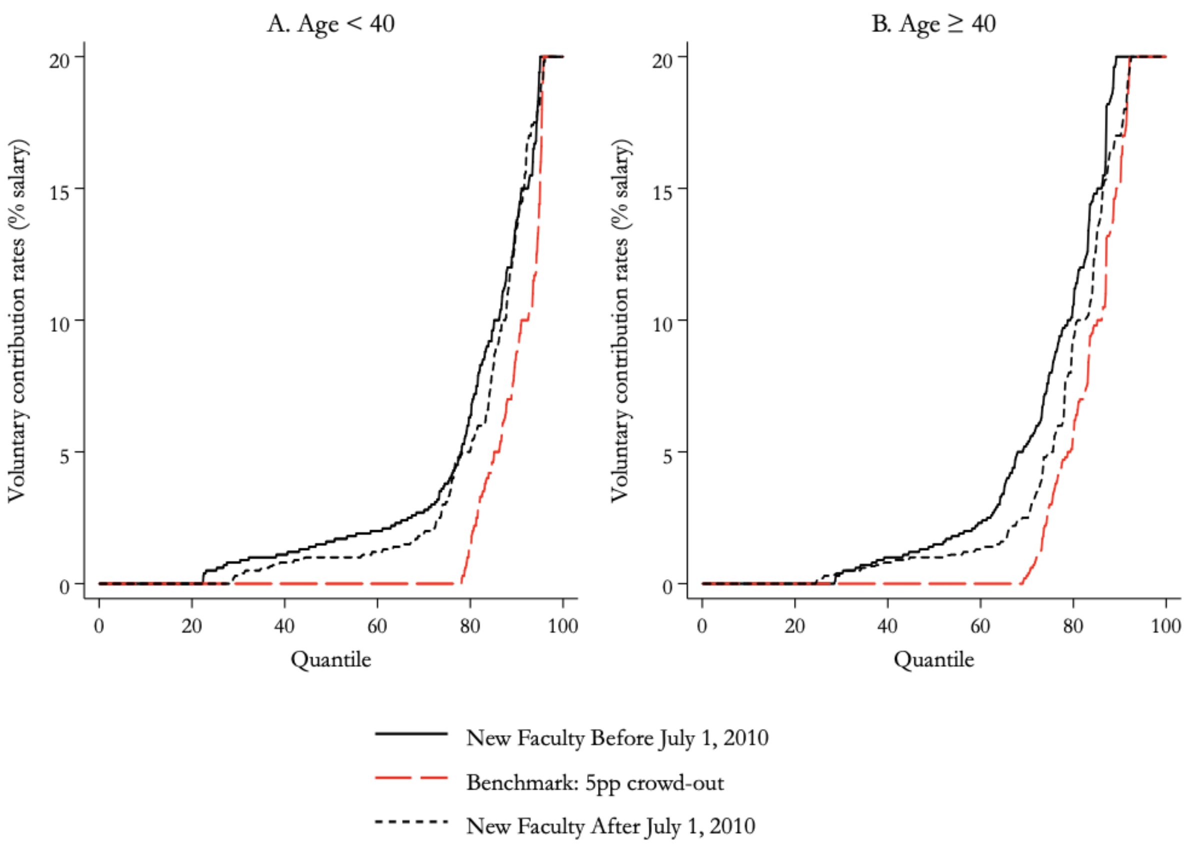

Notes: Figure plots the ordered values of voluntary contribution rates in the 9th month of tenure against quantiles, separately for faculty hired before July 1, 2010 and faculty and faculty hired after this date. Sample excludes faculty who choose the DB plan.

Figure A.11: Quantile regressions of voluntary contribution rates

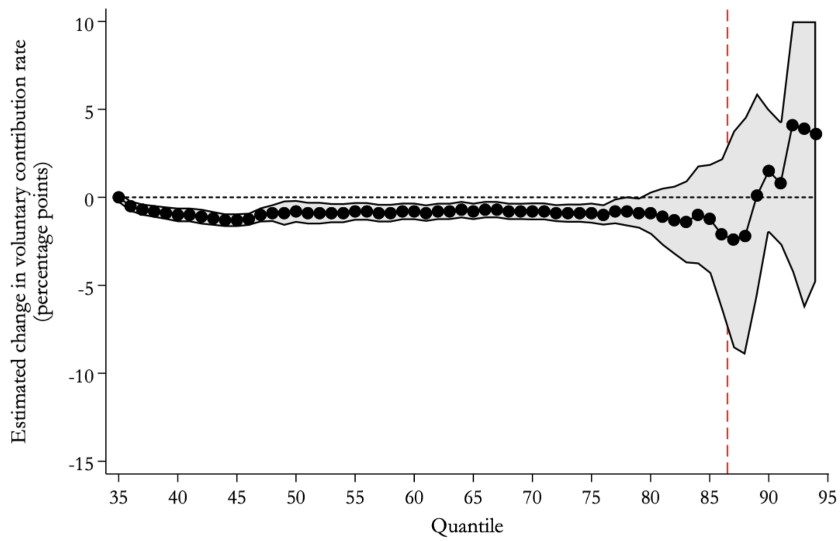

(a) Salary < \$75,000

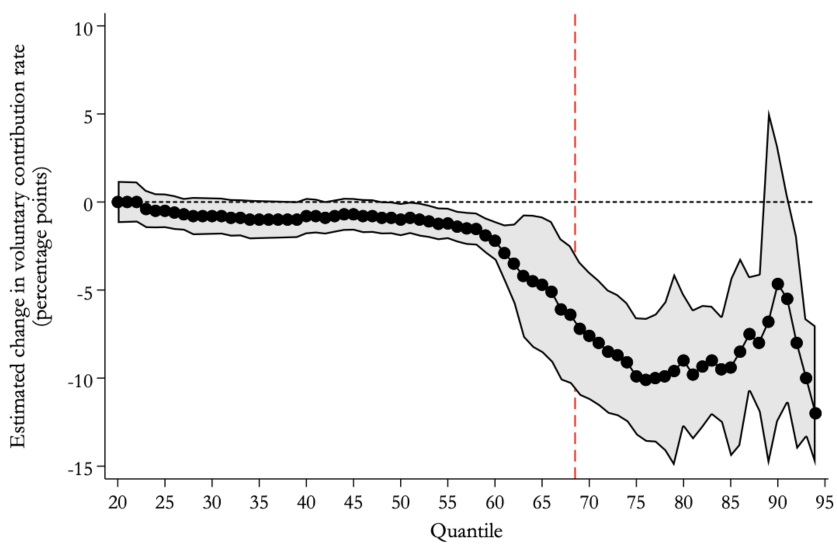

(b) Salary ≥ \$75,000

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.

Consumption Equivalents: Back-of-the-Envelope Calculations This section relates our regression estimates of crowd-out to consumption equivalents, which measure the utility cost or benefits from mandatory saving in dollar terms. To do so, we interpret our estimates of partial crowd-out in response to mandatory contributions using the two-period model from Section 3, which predicts full crowd-out except for those whose liquidity constraints become more binding. A richer model of life-cycle savings would include many more periods and incorporate additional sources of uncertainty. We therefore view the consumption equivalents presented below as a back-of-the-envelope calculation rather than a complete welfare analysis.

While it may be possible that those who stop or were not saving in the first place are harmed and those who save more are helped, in reality we do not know optimal savings for the employees whom we observe. Therefore, we build in flexibility to recognize two pieces of information. We incorporate $\beta \leq 0$ to measure crowd-out, using our estimates from Section 6. We also incorporate a parameter, $\theta \in [0, 1]$, which we vary to calculate consumption equivalents under different assumptions about optimal saving in the absence of the new mandatory required contributions. $\theta = 0$ reflects the case that people were optimizing previously, in which case passive savers who exhibit little crowd-out save too much after the policy is put in place. $\theta = 1$ reflects the case that people were under-saving by the amount of the required contributions. We begin by considering the most straightforward possibilities.

Case 1: Employees make optimal saving decisions. Suppose that employees were already optimizing prior to the policy. In that case, mandatory contributions (weakly) reduce welfare: people are made worse off if they cannot or do not reduce voluntary contributions, or their welfare is unchanged if they fully offset the increase with lower voluntary contributions. We measure this loss from a consumption-equivalent approach using the two-period model from Section 3. 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 lifetime consumption into the future. The consumption equivalent is the amount of money x that solves the following equation:

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

where M_f denotes the mandatory contribution in dollar terms, c_1^* and c_2^* denote optimal consumption in each period, and β denotes the amount of supplemental contributions that are crowded out (with $\beta < 0$ representing crowd-out). Under the assumption that employees were already optimizing, c_1^* equals the employee's salary less voluntary contributions for employees hired before the policy change: $c_1^* = y - M_v^i$, with M_v^i denoting voluntary contributions. Plugging in for c_1^*, the equation becomes:

$$u((y - M_v^i)(1-x)) + \delta u(c_2^*(1-x)) = u(y - M_v^i - M_f + \beta) + \delta u(c_2^* + (1+r)(M_f + \beta)) \tag{A.7}$$

If there is full crowd-out, then $M_f = -\beta$ and there is no welfare loss ($x = 0$) because people fully offset the mandatory increase. If there is incomplete crowd-out, then $M_f > -\beta$ and $x > 0$. People consume "too little" in the first period because $y - M_v^i > y - M_v^i - M_f + \beta$.

They consume "too much" in the second period by the addition of the new savings that are not crowded out, equal to $(1 + r)(M_f + \beta)$. We specify the functional form of utility below when discussing implementation, which allows us to calculate c_2^* from the Euler equation: $u'(c_1^*) = \delta u'(c_2^*)$.

Case 2: Employees are saving optimally after the policy. Now consider the other extreme that people were instead under-saving prior to the policy, and suppose the policy brings them exactly to their optimal savings level. Then the analog to equation (1) is:

$$u((y - M_f - M_v^i + \beta)(1 - x)) + \delta u(c_2^*(1 - x)) = u(y - M_v^i) + \delta u(c_2^* - (1 + r)(M_f + \beta)) \tag{A.8}$$

The left-hand side represents optimal saving and the right-hand side represents the consumption profile in the absence of the policy.

Nested formulation. Nothing in our data directly informs whether Case 1 or Case 2 is more likely, so we introduce a parameter $\theta \in [0, 1]$ that scales the extent to which Case 1 ($\theta = 0$) or Case 2 ($\theta = 1$) applies.[33] θ can also capture the possibility that people crowd out mandatory contributions even if they are undersaving. We set $\theta = 0$ to reflect the assumption that people were already optimizing (Case 1) and $\theta = 1$ to reflect the assumption that they were under-saving by the amount equal to the mandated contribution (Case 2). An intermediate value of θ reflects a situation in which optimal saving lies in between these two extremes. We can then combine equations (2) and (3) by specifying them as a weighted average:

$$\begin{aligned} &u(((1 - \theta)(y - M_v^i) + \theta(y - M_f - M_v^i + \beta))(1 - x)) + \delta u(c_2^*(1 - x)) \\ &= u((1-\theta)(y-M_v^i)+\theta(y-M_f-M_v^i+\beta))+\delta u(c_2^*+(1-\theta)(1+r)(M_f+\beta)-\theta(1+r)(M_f+\beta)) \end{aligned} \tag{A.9}$$

After combining terms, the equation condenses to:

$$\begin{aligned} &u((y - M_v^i - \theta(M_f - \beta))(1 - x)) + \delta u(c_2^*(1 - x)) \\ &\qquad = u((y - M_v^i - \theta(M_f - \beta)) + \delta u(c_2^* + (1 + r)(1 - 2\theta)(M_f + \beta)) \end{aligned} \tag{A.10}$$

This equation nests equations (2) and (3) as special cases by setting $\theta = 0$ or $\theta = 1$, respectively. This setup normalizes the consumption equivalent to equal zero when $\theta = 0.5$. This can be interpreted as it being equally likely the person was already optimizing and is made worse off versus the person was undersaving and made better off.

To implement this approach, we parameterize utility as constant relative risk aversion, with $u(c) = \frac{c^{1-\eta}}{1-\eta}$ and $\eta = 3$ denoting the coefficient of relative risk aversion. We assume an annualized discount factor of $\delta = 0.98$ and real interest rate of 2 percent. While the equations above use a two-period model, the two periods are not in consecutive years but rather while working and in retirement. We assume the periods are 15 years apart and adjust the discount factor and interest rate accordingly. We then calibrate the rest of the equation

[33] If people were over-saving in the absence of the policy, then the welfare losses will be larger than those from equation (2).

using data of y and M_v^i from the empirical distributions shown in Figure 5, calculating M_f as 5 percent of y, and using the estimates of β for each salary quantile using the quantile regression results from Appendix Figure A.11.

Figure A.12 plots the results of solving equation A.10 to calculate consumption equivalents for 20 evenly-spaced values of θ ranging from 0 to 1. For each employee in the pre-period, we solve for the consumption equivalent in percentage terms, and then convert to dollars based on their salary, voluntary contributions, and estimated crowd-out rate β from the quantile regressions. We then average over everyone in the sample to construct the overall consumption equivalent (solid line). To explore differences in welfare effects by salary level, we also calculate averages by employees earning below or above $75,000.

If $\theta = 0.5$, half the people are helped and half are hurt by the new policy, while more are hurt (the average consumption equivalent is negative) for smaller values of θ. If people were already saving optimally ($\theta = 0$), then partial crowd-out by both active savers who stop participating (but cannot reduce saving by more) and lack of crowd-out by passive savers (who should not save more) brings the average welfare loss to about $1,765 per capita in each period. In percentage terms, this amounts to a 3.0% loss in consumption. These costs vary by income level. The welfare losses are about $2,155 for lower earners and $1,415 for higher earners. Although higher earners have a greater willingness-to-pay for a smooth consumption profile, they exhibit substantial crowd-out and, as a result, are less affected by the policy. Low earners, by contrast, exhibit less crowd-out and a larger fraction have no voluntary contributions to begin with. If both groups were already saving optimally, the increase in mandatory saving hurts lower earners more.

If we instead assume that people were initially undersaving and the policy brings them to their optimum ($\theta = 1$), then the average benefits are about $2,225 (equal to a 4.0% gain in consumption). The benefits are now larger for lower earners, again because they exhibit less crowd-out than higher earners. At intermediate values of θ—which weigh the extent to which a person was already saving optimally vs. undersaving—the consumption equivalents are lower. By construction, there is no benefit or cost when $\theta = 0.5$, which represents an even split between the two extreme cases.

Figure A.12: Consumption equivalents by θ and Salary

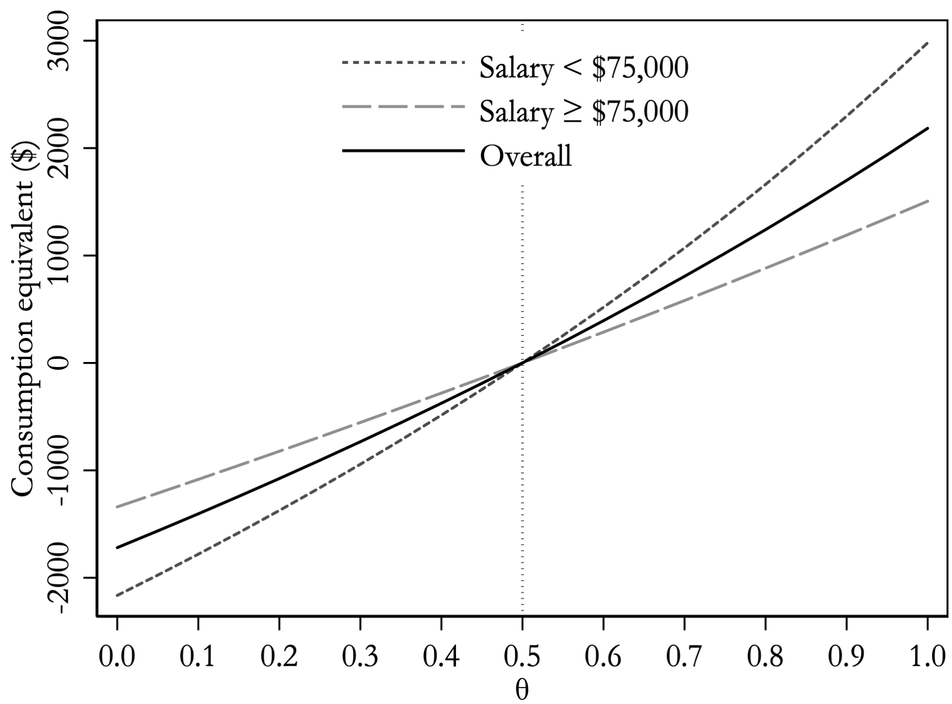

Notes: Figure plots the consumption equivalents from solving equation A.10 as a function of θ, which governs the degree to which employees were saving optimally in the absence of the policy. $\theta = 0$ represents the case that people were already saving optimally without the mandatory contributions, and so the policy unambiguously reduces utility. $\theta = 1$ represents the case that mandatory saving brings people exactly equal to their optimal saving level. By construction, there are no costs or benefits when $\theta = 0.5$ as the two forces exactly offset at this level.

Internet Appendix 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.