

Smaller than We Thought? The Effect of Automatic Savings Policies

James J. Choi
Yale University and NBER

David Laibson
Harvard University and NBER

Jordan Cammarota
University of California, Berkeley

Richard Lombardo
Harvard University

John Beshears
Harvard University and NBER

December 17, 2024

Abstract: Medium- and long-run dynamics undermine the effect of automatic enrollment and default savings-rate auto-escalation on retirement savings. Our analysis of 401(k) plans incorporates the facts that employees frequently leave firms (often before matching contributions have fully vested), a large percentage of balances are withdrawn upon employment separation, and many employees opt out of default auto-escalation. Across nine natural experiments, steady-state saving rates increase by 0.6% of income due to automatic enrollment and 0.2% of income due to default auto-escalation. Only 43% of employees in 21 401(k) plans in our data with default auto-escalation actually escalate on their first auto-escalation date.

Contributions: conceptualized paper (JJC, DL, JB), wrote paper (JJC), edited paper (DL, JB), implemented empirical analysis (JC, RL), project leadership (JJC, DL, JB). We thank Jessica Brooks, Harry Kosowsky, Justin Katz, and Lea Nagel for excellent research assistance, and Shlomo Benartzi, Richard Thaler, and audiences at LBS, Northeastern, Princeton, University of Southern California, Stanford, and University of Washington for useful comments. The research reported herein was performed pursuant to grant RDR18000003 from the US Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium, and received additional support from the Eric M. Mindich Fund for Research on the Foundations of Human Behavior and the Pershing Square Fund for Research on the Foundations of Human Behavior. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA, any agency of the Federal Government, or NBER. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation, or favoring by the United States Government or any agency thereof.

Automatic enrollment (where an employee contributes to a savings plan unless they opt out) and default auto-escalation (where an employee's saving rate in the plan increases automatically over time unless they opt out) have become major parts of retirement savings policy around the world.¹ Forty percent of U.S. private industry workers participating in a savings and thrift plan do so in one with automatic enrollment (Zook, 2023), and 40% of plans administered by Vanguard automatically escalate employee contributions unless employees opt out (Vanguard, 2023). The SECURE 2.0 Act requires most 401(k) retirement savings plans established after December 29, 2022, to automatically enroll eligible employees into the plan and auto-escalate their savings rate unless they opt out.² Ten U.S. states have passed legislation requiring employers that do not sponsor a 401(k) plan to automatically enroll employees in an Individual Retirement Account.

In this paper, we estimate the effects of such automatic policies on retirement savings accumulation, accounting for medium- and long-run factors that have not been jointly evaluated in prior work. Our results indicate that these medium- and long-run factors meaningfully reduce the impact of automatic policies on accumulation in the U.S. retirement savings system.

Much research has found that automatic enrollment greatly increases the fraction of employees who contribute to the 401(k) and modestly increases average contribution rates (Madrian and Shea, 2001; Choi et al., 2002, 2004; Beshears et al., 2008; Derby, Mackie, and Mortenson, 2023). Choukhmane (2024) uses an estimated structural model to conclude that automatic enrollment only slightly increases average retirement savings over the lifecycle.³ Thaler and Benartzi (2004) find that employees who opt into auto-escalation experience repeated annual increases in their contribution rates, and Benartzi, Peleg, and Thaler (2013) report that the take-up rate of auto-escalation increases from about 25% to about 85% when it is made the default.

¹ Italy, Lithuania, New Zealand, Poland, Turkey, and the United Kingdom have implemented automatic enrollment at the national level (OECD, 2021).

² This requirement takes effect for plan years beginning January 1, 2025.

³ His analysis focuses on a benchmark case in which a 50% employer matching contribution that vests slowly over two years, empirically calibrated levels of job turnover, illiquid defined contribution (DC) retirement accounts, and a \$250 switching/transaction cost (incurred when opting in/out of a DC plan or actively changing a DC contribution rate) jointly explain why (1) newly hired workers are not strongly incentivized to contribute to their 401(k) plan and hence do not quickly opt in, (2) auto-enrolled workers do not opt out, and (3) non-autoenrolled workers opt in eventually as their vesting status improves. In the model, withdrawals can only be made from the DC account when the worker is unemployed, and not when the worker is employed or when the worker is switching between jobs without an official unemployment spell. Choukhmane (2024) provides extensive robustness analysis, which explores the ways that the model's predictions change as the 401(k) plan parameters change, including different match rates and vesting schedules.

Benartzi and Thaler (2013) estimate that auto-escalation increased U.S. retirement savings by \$7.4 billion in 2013 under the assumption that auto-escalation increases the savings rate of those who accept it by 3% of income relative to the counterfactual without auto-escalation.

In our main analysis, we study nine firms that, sometime between 2005 and 2011, introduced either (1) automatic enrollment on its own, (2) default auto-escalation in a 401(k) plan where automatic enrollment was already present, or (3) automatic enrollment and default auto-escalation simultaneously. The automatic policies applied only to employees hired from a certain date onward, so we identify their effect by comparing 62,430 employees hired in the year after the policy introductions to 55,937 employees hired in the year before the policy introductions. We incorporate three methodological advances that have never been simultaneously present in prior studies.

First, we estimate long-run treatment effects averaging over both high-turnover and low-turnover employees, assuming departing employees move to a firm with the same 401(k) structure and repeat the contribution and turnover behavior they chose at their original firm. Most previous analyses of automatic policies estimate a treatment effect at time t conditional on remaining at the original firm at t . This approach makes sense for studying the narrow question of how continuous treatment under an automatic policy affects outcomes for non-separating employees (over horizon t). However, in practice, many new hires leave the firm after only a short time. In 2022, about 4% of the U.S. nonfarm labor force left their current job each month, according to the Job Openings and Labor Force Turnover Survey. In the universe of Vanguard-administered savings plans with an automatic savings feature, 29%, 42%, and 52% of new employees have left the firm by one, two, and three years after hire, respectively (Clark and Young, 2021). In principle, we would like to estimate the population-level average impact of an automatic policy that is permanently implemented for the entire population, keeping in mind that workers who depart before the automatic enrollment opt-out deadline or the first automatic escalation date are unaffected by these automatic policies, and workers who move between auto-escalation firms are restarted at the lowest contribution rate in the escalation ladder. Furthermore, leaving a firm early has a direct negative impact on retirement savings accumulation if it causes one to forfeit accrued employer 401(k) matching contributions due to not having worked for the firm long enough for those employer contributions to be fully vested.

Second, we account for the fact that when employees separate from a firm, they withdraw a large percentage of their 401(k) balances, either voluntarily or because employers are allowed by law to compel a complete liquidation of account balances below \$1,000 (unless the employee provides instructions to roll over the balances to an Individual Retirement Account). Argento, Bryant, and Sabelhaus (2015) report that for every dollar contributed to U.S. defined contribution retirement accounts in 2010 by those under the age of 55, approximately 40 cents leaked out to this age group as a pre-retirement withdrawal in the same year. Because automatic savings policies have the strongest positive contribution effect on those with weak savings motives, much of the asset accumulation induced by automatic policies during the employment spell may be withdrawn upon separation.

Third, we credit auto-escalation only for contribution increases that actually occur. Except for one 401(k) plan in their sample, Benartzi, Peleg, and Thaler (2013) are only able to measure opt-outs from default auto-escalation by observing whether employees have opted out *before* their first scheduled contribution increase. They mostly do not observe whether employees' contribution rates indeed increase on the auto-escalation date.⁶

To eliminate noise induced by the cohort-specific realization of asset returns, our main outcome variable is based on how much individuals would have accumulated in the 401(k) system during their tenure at the observed job if their 401(k) portfolio returns had been a constant 5% per year. To implement this calculation, we use each employee's longitudinal contribution rates and employment duration, along with the population-level⁷ empirical relationship between 401(k) balances at separation and cash (i.e., non-rollover) withdrawals. We convert this net-of-cash-distributions accumulation amount (which includes employer contributions that are actually received because they are vested) into the equivalent constant contribution rate that would have resulted in the same net accumulation by the end of the employee's tenure. In summary, our equivalent constant contribution rate is the (smoothed) annualized individual 401(k) savings rate

⁶ Benartzi, Peleg, and Thaler (2013, p. 250) write, "One caveat is that the opt-out program was generally introduced in 2005 with the first saving increase scheduled for 2006. Hence, we could not determine from our data, which ends with 2005, how many participants, if any, opted out right before the increase. Data from the one plan that introduced the program in 2004 and already had the first increase in 2005 suggest an opt-out rate of just 9%, so it does not look like participants opted out right before the first increase."

⁷ This is calculated from a sample of 7,412 employees at eight firms, which is smaller than the entire sample for data reasons we discuss later.

that *would* have produced the same amount of net wealth formation that we calculate for each employee.

Because the equivalent constant contribution rate annualizes each employee's net savings rate inclusive of employer contributions over their (single-employer) tenure, it is comparable across employees who remain at the company for different lengths of time. The average of these equivalent constant contribution rates across employees in a cohort is the steady-state annualized cohort-average contribution rate under the assumption that separating employees repeat their contribution, withdrawal, and attrition behavior in every subsequent job spell, and their subsequent employers have the same 401(k) features as the employer we observe in our data. The difference in this average between treatment and control cohorts is the steady-state treatment effect on annualized net retirement savings rates.

Table 1 summarizes the treatment effects of the three auto policies, averaged over the firms that implemented each policy and estimated using several different methodologies to compute the equivalent constant contribution rate. The first row extrapolates treatment effects at firms with auto-escalation using only contribution rates from the first year of tenure, ignoring the wealth-depleting effects of cash distributions and incomplete employer-match vesting and assuming 85% take-up of auto-escalation; this set of assumptions is a natural way to extrapolate from the results of early pioneering papers in this literature (Madrian and Shea, 2001; Benartzi, Peleg, and Thaler, 2013). The second row adjusts the calculations presented in the first row by using observed take-up rates of auto-escalation on the first date on which contribution rates were due to escalate by default. This row also shows the automatic enrollment treatment effect extrapolated from the first year of tenure, still ignoring the wealth-depleting effects of cash distributions and incomplete vesting. The third row repeats the analysis of the second row, but uses data from five years of tenure, which is similar to the way long-run effects were estimated by Choi et al. (2002, 2004), Thaler and Benartzi (2004), and Beshears et al. (2008). The fourth row incorporates the wealth-depleting effect of incomplete vesting of employer contributions when employees leave the firm before becoming fully vested. The fifth row additionally incorporates the effect of 401(k) cash distributions; this row is our most comprehensive estimate.

Relative to the first treatment effects in each column, the average (equally weighted across all nine companies at which we study a natural experiment) most-comprehensive treatment effect

on savings rates is smaller by 77%: 0.6 percentage points of income compared to 2.6 percentage points of income.

We next discuss the treatment effects for each specific type of natural experiment. For our most comprehensive treatment effects, the average increase in the equivalent constant contribution rate across the four companies that introduced *automatic enrollment alone* is 0.6 percentage points of income. The average increase in the equivalent constant contribution rate across the two companies that introduced *default auto-escalation on top of pre-existing automatic enrollment* is 0.2 percentage points of income. The average increase in the equivalent constant contribution rate across the three companies that *simultaneously introduced both automatic enrollment and default auto-escalation* is 0.8 percentage points of income.

The savings effects are modest for four primary reasons. First, many individuals *not* subject to automatic policies actively choose contribution rates over time that erode much of the savings gap that would exist if individuals were completely passive. Employees not subject to automatic enrollment increase their contribution rates more quickly than automatically enrolled employees, and most of the contribution increases that auto-escalation implements occur even in the absence of auto-escalation.

Second, opting out of auto-escalation is surprisingly common when it is the default. Among employees who are still working at our five default auto-escalation firms on their first escalation date, only 40% escalate, and acceptance of the escalation default decays further at each subsequent escalation date. Figure 1 shows auto-escalation acceptance rates in a larger sample of 21 default auto-escalation firms (discussed in Appendix B). On average, the acceptance rate of the auto-escalation default is 43% on the first escalation date, 36% on the second date, and 29% on the third date.⁸ In our analysis, being counted as “accepting” auto-escalation on the second and third escalation dates does not require acceptance on all previous escalation dates. In other words, an employee who opts out of auto-escalation and then later opts back in is counted as accepting auto-escalation on all escalation dates after they opted back in.

Third, the employee turnover rate is high, and many employees leave before their employer matching contributions are fully vested, resulting in significant forfeited money.

⁸ The third escalation date’s acceptance rate is averaged over 20 firms, as one firm is missing contribution rate data for the third date.

Fourth, 42% of 401(k) balances are cashed out upon separation from an employer. Controlling for balance at separation, there is little difference in cash leakage rates between those who are and are not subject to automatic policies. However, individuals in treated cohorts are more likely to separate with relatively small (but non-zero) 401(k) balances, so the average leakage rate as a fraction of 401(k) balance at separation among employees with positive balances is 9 percentage points higher in treated cohorts than in control cohorts.

Our sample of firms is small, which means that our estimated net effect of automatic policies could differ substantially from an estimate in a larger, more representative sample. Nevertheless, statistics from other large samples suggest that the erosion of automatic policy treatment effects through the additional channels we consider is also sizable in the general population. Alight (2019) reports that in their universe of more than 2 million participants, 40% of departing workers cash out their entire 401(k) balance. The treatment effect of automatic policies would shrink by 40% relative to an estimate that failed to consider withdrawals if (a randomly chosen) 40% of both treated and untreated employees liquidate their 401(k). Carranza and Goodman (2024) find that in the Vanguard universe, 30% of employment separations occur before the employer match is fully vested, and among those forfeiting match money, 40% of total account balances is forfeited on average. Thus, the treatment effect of automatic policies would shrink by 12% due to forfeitures in the absence of withdrawals if treated and untreated employees forfeited matches at the same rate. In a sample of 46 plans, Choukhmane (2024) estimates that among workers who do not leave the firm, automatic enrollment ceases to increase cumulative contributions after the first year of tenure because untreated employees choose higher contribution rates over time. Zhong (2021) finds that in the OregonSaves auto-IRA program, the acceptance rate of the auto-escalation default on the first escalation date among previously passive savers is only 47%.

Clark and Young (2021) report that in the Vanguard universe, the acceptance rate of an auto-escalation default is 63%, 63%, and 60% after one, two and three years of tenure—significantly below the 85% found by Benartzi, Peleg, and Thaler (2013), but significantly higher than in our sample.⁹ Our sample's lower acceptance rate may be due to higher salience of each

⁹ In the U.K., there has been no significant increase in opt-out rates as the minimum allowable default pension contribution rate, which is also the minimum allowable pension contribution rate, increased from 2% to 5% to 8% of income (Department for Work & Pensions, 2022). However, opting out also became increasingly costly, as the resulting foregone employer contribution rose from 1% to 2% to 3% of income.

auto-escalation date for employees in our sample compared to employees in the Vanguard sample.¹⁰

In the final row of Table 1, we report average treatment effects when we randomly assign additional treatment employees to accept auto-escalation such that we match the higher Vanguard take-up rates. We optimistically assume that none of the contribution rate increases implemented for these randomly chosen employees would have occurred without automatic escalation. As expected, the treatment effects increase—to 0.9% of income for auto-escalation alone and 1.1% of income for simultaneous auto-escalation and autoenrollment. These numbers remain far smaller than auto-feature savings effect estimates that ignore endogenous contribution rate catch-up, job turnover, incomplete vesting, and withdrawals. Averaged across the five firms that introduced auto-escalation (either alone or in conjunction with automatic enrollment), raising the assumed auto-escalation take-up rate increases the treatment effect relative to the comprehensive estimate by only 0.4% of income.

Our results suggest that the moment of job transition is a key weakness in the U.S. retirement savings system. Beshears, Choi, Laibson, and Maxted (2022) present a model where the same present bias that causes automatic savings policies to initially increase savings also dissipates those savings when households are allowed to withdraw them after job separation. Making retirement savings accumulated at previous jobs less easily accessible, creating infrastructure that allows high contribution rates attained at a previous job to be automatically enacted at one's next job, and/or making the auto-escalation contribution level dependent on something like employee age instead of only tenure at one's current job may significantly increase the impact of automatic policies on retirement savings accumulation. Beshears et al. (2024) find in a two-period setting with unobserved taste shocks that if the strength of present bias is sufficiently heterogeneous across the population, the optimal savings system makes a large fraction of net worth—almost enough to smooth consumption across the retirement transition on its own—*completely* illiquid before retirement.

¹⁰ Employees in our sample who actively change their contribution rate do not appear to be automatically kicked off of auto-escalation, judging from the sizable fraction of employees who increase their contribution rate on escalation dates but are not on the completely passive contribution path. From personal communication, we have learned that two other 401(k) plan administrators, Voya and Empower, do automatically stop auto-escalation for active contribution changers in plans they administer, while at Vanguard plans, contribution rate changers must choose whether they want to auto-escalate on a screen where their previous auto-escalation election is pre-selected.

Our paper is related to others that investigate whether savings nudges are offset at un-nudged margins. Most closely related is Derby, Mackie, and Mortenson (2023), who estimate in a large sample of firms that some of automatic enrollment’s savings effect is undone by subsequent withdrawals, but they find there is no offset via reduced retirement contributions by the employee’s spouse. We are distinguished from Derby, Mackie, and Mortenson (2023) in that we analyze auto-escalation in addition to automatic enrollment; we are able to observe employer contributions, which are an important component of 401(k) wealth accumulation; and we estimate longer-run treatment effects using cohorts hired close together in time, whereas Derby, Mackie, and Mortenson (2023) need to compare cohorts hired five years apart from each other, heightening concerns about confounding calendar time effects.

Choukhmane (2024) finds that automatic enrollment in the current job’s pension causes workers to become less likely to save in their next job’s pension if they must opt into it in order to participate, but there is no such dynamic offset if the next pension also uses automatic enrollment. Beshears, Choi, Laibson, Madrian, and Skimmyhorn (2022) report no statistically significant effect of automatic enrollment on financial distress, credit scores, or debt excluding auto and mortgage debt. Using a much larger sample, Beshears et al. (2024) find that about 20% of the pension savings created by automatic enrollment is offset by increased unsecured debt—an estimate that is within the 95% confidence interval of the corresponding Beshears, Choi, Laibson, Madrian, and Skimmyhorn (2022) estimate—and automatically enrolled individuals become more likely to have a mortgage but less likely to go into financial distress. Blumenstock, Callen, and Ghani (2018) find that automatic enrollment does not reduce other savings, although their estimates are imprecise. Chetty et al. (2014) estimate that 30% of increases in compulsory pension savings are offset by reductions in non-pension savings and increases in debt.

Outside of the retirement saving context, Medina (2021) shows that credit card payment reminders increase checking account overdraft fees, while Guttman-Kenney et al. (2023) find that shrouding the option to automatically make only the minimum monthly credit card payment has no effect on debt reduction because of offsetting consumer responses. Medina and Pagel (2023) find that a text message encouraging savings is successful at increasing saving among a subset of recipients while not increasing their borrowing. Brown, Grozicki, and Medina (2023) report that limiting the marketing of credit cards to college students increases their student loan balances.

More generally, DellaVigna and Linos (2022) find that nudges (excluding those that change the default) implemented by government nudge units have smaller effect sizes than those studied in published academic studies, with 60-70% of the shrinkage being attributable to publication bias. Jachimowicz et al. (2019) perform a meta-analysis of default effects that finds that defaults have a large impact on average. Altmann, Grunewald, and Radbruch (2024) use evidence from a laboratory experiment to argue that good defaults have positive spillover effects to other domains because they free up cognitive resources that can be devoted to un-nudged tasks.

The paper proceeds as follows. Section I describes how we select our sample of firms and the nature of our data. Section II discusses how we choose our control and treatment cohorts and construct the equivalent constant contribution rate outcome variable. Section III describes the results of our estimations. Section IV concludes. Appendix A shows that cash leakage from 401(k) accounts tends to occur shortly after job separation, which justifies the way we treat cash leakage in constructing our outcome variable. Appendix B describes our results on auto-escalation acceptance in a sample of 21 firms, which is larger than the sample of five firms with default auto-escalation used in our main text's analysis.

I. Firm selection and data description

Our 401(k) administrative data come from Alight, a company whose services include providing defined contribution pension recordkeeping services for employers. Within a universe of approximately 200 firms, we identified 86 instances of automatic savings policies that were implemented between January 1, 2003 and January 1, 2011.¹¹ We then imposed the following requirements: (1) the policy only affected employees hired after the policy was introduced; (2) additional automatic savings policies were not introduced during the time period over which we measure outcomes, which extends five years after an employee cohort is hired; and (3) data on contribution rate elections were available during the study period. These conditions restricted our sample to nine automatic policy implementations.

¹¹ We identified the introduction of automatic savings policies from a survey of Alight clients conducted in 2010 and 2019, by reading plan documents that span a range of years that differs for each company, and by searching for large, discontinuous, and sustained increases in the number of employees at a certain contribution rate, which is indicative of a new contribution rate default or auto-escalation. If there are automatic savings policies that we failed to identify from the administrative data because their effect was more muted than what the literature has previously documented, then our sample is biased towards more "successful" policies.

Table 2 describes the 401(k) plan features at each firm in our sample. We divide our firms into those that introduced automatic enrollment only (“autoenrollment firms”), those that introduced default auto-escalation in a context where they were already automatically enrolling employees (“auto-escalation firms”), and those that introduced automatic enrollment and default auto-escalation simultaneously (“autoenrollment and escalation firms”). The initial default contribution rates range from 2% to 6% of income. At firms that introduce default auto-escalation, contribution rates automatically increase by 1% of income per year until a maximum that varies from 6% to 15% of income across firms. At Firms E through H, the first auto-escalation date did not occur within the first year of tenure for some employees. At Firm E, even though employees hired from January 2011 onwards were subject to auto-escalation every April, the first escalation date was not until April 2012. At Firms F, G, and H, auto-escalation occurs every January, but only for employees with at least six, five, and nine months of tenure, respectively, so employees hired later in the calendar year do not automatically escalate until their second year of tenure. Firm I is the only firm that escalates on the anniversary of the employee’s hire.¹²

All of our firms match employee contributions, and Firm I makes additional employer contributions that do not depend on the employee’s contribution choices starting in April 2009. Two of the autoenrollment and escalation firms increased the generosity of their match rates in the middle of the study period, and one increased its match threshold (the savings rate up to which contributions are matched) for a small number of its employees. Three firms feature immediate 100% vesting of employer contributions—employees forfeit none of their employer contributions no matter when they leave the company—but most companies require a certain length of tenure before employer contributions are fully vested. Four firms have cliff vesting, where employees jump from having 0% to having 100% of their contributions vested upon exceeding three years of tenure. Two other firms gradually increase their vesting percentage as tenure increases up until five years. This distribution of vesting schedules is somewhat less generous than what Vanguard (2013) reports for its universe in 2012, where 44% of plans immediately vest, 18% of plans have a five-year graded vesting schedule, and 10% have three-year cliff vesting.

We have two types of administrative data. The first data set is a series of cross-sections at year-end for each firm in our study. Each cross-section contains employee-level information on

¹² More exactly, it escalates either zero, one, or two months after the employee’s hire anniversary, depending on the employee.

birth date, hire date, gender, salary, and job separation date. It also contains year-end 401(k) plan balances; total dollars contributed to the plan (separately for employer and employee contributions) during the year; and for each withdrawal, the date of the transaction, the total dollars withdrawn, and the total dollars rolled over to an outside retirement account. The second data set contains monthly 401(k) contribution rate elections for each employee. These contribution rates are chosen as a fraction of salary (rather than a fixed dollar amount) and can be changed by employees at any time.

II. Methodology

A. Cohort construction

We identify the effect of automatic policies by comparing two hire cohorts within each company. The treatment cohort was hired within the 365 days after the policy introduction date, and hence was subject to an automatic policy. The control cohort was hired within the 365 days prior to the policy introduction date, and hence was never subject to the automatic policy. This is the empirical approach used by many prior studies of automatic enrollment (Madrian and Shea, 2001; Choi et al., 2002, 2004; Beshears et al., 2008; Choukhmane, 2024).¹³ We exclude employees who were re-hired (because it is not clear whether they would be subject to the automatic policy) and employees who ever rolled money from another plan into the plan we observe (because their 401(k) balance at employment separation would not correspond closely to what is implied by their contribution rate path during employment).

Table 3 shows the characteristics of the treatment and control cohorts at each firm. As expected for new hires, the average age of the cohorts is relatively young—between 30 and 38. There is a wide range of gender compositions, from 19% to 78% female. Average starting salary runs from \$23,396 to \$64,715. As a point of comparison, Clark and Young (2021) report that in the Vanguard universe as of June 2020, employees hired between 2017 and 2019 into a firm whose 401(k) has an automatic savings feature are 41% female and have a median income and age of \$40,300 and 33. A large fraction of our sample’s employees has left the firm within the first five

¹³ It is in principle possible that this methodology underestimates the treatment effect because the control cohort increases its contribution rate in response to the introduction of the automatic policy. We informally test for such an effect by comparing the first six months of contribution rates for control cohort individuals hired 186 to 365 days before the policy implementation date to contribution rates for control cohort individuals hired 185 to 1 days before. In only two of the seven firms that implemented automatic enrollment is there a hint that later-hired control individuals contribute more.

years after hire—between 43% and 90%—highlighting the importance of incorporating the effects of attrition when estimating the impact of automatic savings policies.

Statistical tests indicate that the above means and proportions are often significantly different between cohorts within a firm. We will control for observable differences between cohorts in a regression framework. Whether or not our treatment effect estimates are biased by differences in unobservable characteristics, our paper’s main point still stands: contribution rate catch-up, job attrition, incomplete vesting, withdrawals, and low acceptance of auto-escalation undermine the impact of automatic policies on retirement savings accumulation. We have also tried choosing cohort hire windows that minimize the observable differences between treatment and control cohorts.¹⁴ Online Appendix Table 1 shows that this exercise is successful in reducing the imbalances between cohorts, and Online Appendix Table 2 shows that the estimated treatment effects under this approach are qualitatively similar to the ones we present in Section III.

B. Outcome variable construction

Our goal is to estimate the effect of automatic savings policies on retirement wealth accumulation. One way to do this is to compare retirement account balances of the treatment cohort to retirement account balances of the control cohort at equivalent times since hire. However, such a comparison would yield a highly imprecise estimate because each cohort has experienced different cumulative capital gains at each tenure time due to random fluctuations in the capital markets. In addition, the comparison would be confounded by the fact that the introduction of automatic enrollment is accompanied not only by a change in the default contribution rate, but also by an introduction of an asset allocation default that almost always differs from the average asset allocation chosen by participants not enrolled under automatic enrollment. We are seeking to estimate the effect of automatic savings policies purged of the effects of noise in capital gains and differential asset allocations.

¹⁴ Potential treatment cohorts are employees hired in the 90, 180, 270, or 365 days after the automatic policy was implemented. Potential control cohorts are employees hired in the 90, 180, 270, or 365 day window that either begins one year before policy implementation or ends one day before policy implementation. We only consider cohort pairs where the treatment and control cohorts have the same hire window width, which means that we have seven candidate treatment-control cohort pairs for each firm. We choose the cohort pair within each firm that has the smallest equal-weighted average of three differences between the cohorts: mean deflated salary divided by its standard deviation (pooled across both cohorts), mean age divided by its pooled standard deviation, and percent female.

Therefore, we instead choose to analyze as our outcome variable a synthetic measure of retirement wealth accumulation achieved under a hypothetical scenario where employees' investment returns are a constant 5% per year and employees' contribution rate path equals what we actually observe for them.¹⁵ Let this balance at the end of employee i 's tenure, inclusive of employer contributions but subtracting off post-separation withdrawals, be b_{iT} . We map b_{iT} for each employee to their "equivalent constant contribution rate." The equivalent constant contribution rate is the constant employee contribution rate during employment (with no withdrawals eroding accumulation) that would result in the employee having b_T at employment separation if annual returns are always 5% and there are no employer contributions. Estimating the treatment effect on this equivalent constant contribution rate allows us to interpret the treatment effect as the annualized increase in savings rates that would result from applying a given automatic policy forever (including at future jobs) instead of applying the control policy forever, under the assumption that an employee who departs a sample firm will have the same 401(k) features and will exactly repeat their contribution, withdrawal, and attrition behaviors in every subsequent job spell.

In this subsection, we describe the construction of the most comprehensive version of the outcome variable. When we discuss our results in Section III, we will present treatment effects of automatic policies not just on this most-comprehensive outcome variable, but also on outcome variables that exclude or modify some of the steps in the outcome's construction.

1. Contribution accumulation

We begin with employee contribution rate elections, which are observed monthly and expressed as a percent of salary. We assume that contributions are deposited at the very end of each month. Because of gaps in our data on salary and unnormalized dollars contributed to the 401(k), our approach cumulates contribution rates over time without adjustment for salary growth.

While individual i is employed at the firm, their synthetic balance b at the end of month t follows the law of motion

¹⁵ Choi et al. (2009) find that 401(k) contribution rates increase in response to personally experienced 401(k) returns that are high on average with low variance, but the magnitude of these effects is small relative to the effects of automatic policies. For example, moving from a 0% to a 1% monthly return over a year causes the 401(k) contribution rate to contemporaneously rise by 0.08% of income. An increase in monthly return standard deviation from 4% to 6%—roughly a movement from the 50th to the 75th percentile in their sample—causes the 401(k) contribution rate to fall by 0.08% of income contemporaneously and another 0.08% of income the following year.

$$b_{it} = 1.05^{1/12}b_{i,t-1} + c_{it}(1 + m_{it}v_i)/12 \quad (1)$$

where c_{it} is i 's contribution rate at t , m_{it} is the percent of i 's contribution at t that is matched by the employer, and v_i is the percent of the employer match that is vested at the end of the employee's tenure at the company.¹⁶ We divide the monthly contribution rate by 12, which makes our synthetic balance variable normalized by annual salary.

Contribution rates are sometimes missing in our data. If they are missing for every month of an employee's tenure and the employee remained at the firm for at least six months, then we drop the employee from the sample. This criterion excludes 1% or fewer of the employees hired in the two-year window centered on the policy introduction date at seven firms, but excludes 3% and 7% of such employees at firms F and G, respectively. For employees who are always missing contribution rates and who leave the firm within six months, we impute zero contribution rates for their entire tenure, under the theory that it is plausible that a very short-tenured employee who never enrolls in the 401(k) would not be entered into the contribution rate database. For the remaining employees, if their missing contribution rates are at the beginning or end of their tenure, we impute the closest non-missing contribution rate to those observations. Otherwise, we impute missing observations using linear interpolation between the last non-missing contribution rate before the missing observations and the first non-missing contribution rate after the missing observations. The bottom of Table 3 shows that we impute 9% or fewer contribution rate observations at all firms except Firm H, which has the greatest percentage of missing observations (25%) because contribution rate data are unavailable for two calendar years during the study period.

2. *Correcting for changes in the employer match structure over time*

At firms H and I, the employer match became significantly more generous in the years after each firm implemented its automatic policy. Because the control cohort was hired prior to the treatment cohort, these match changes occur later in tenure time for the control cohort than for the treatment cohort. In order to approximate what the two cohorts' retirement accumulation would have been had both of them experienced the same path of 401(k) match structures in tenure time,

¹⁶ Firms F, H, and I offer different matches to different employees, but our data do not explicitly identify which match structure an employee has. We infer each employee's match structure from the ratio of employee contribution dollars to employer contribution dollars within a calendar year, combined with the path of contribution rate elections during that calendar year. We assign the most common match structure within the company to those whose dollar contribution data are missing.

we calculate synthetic balances for the control cohort under the assumption that any change in the match structure occurred twelve months earlier in calendar time than its actual date for its members.¹⁷ More generous matches are likely to increase average employee contributions (Choi, 2015), but we do not alter the contribution rates chosen by employees when doing our calculations. This biases us towards estimating a more positive treatment effect at firms H and I.

3. Withdrawals while employed at the 401(k) sponsor

Employees can make withdrawals from their 401(k) under limited circumstances while still employed at the 401(k) sponsor. We estimate the fraction of the individual's 401(k) balance distributed in such an "in-service withdrawal" at month t by dividing the withdrawal amount at t by the sum of the 401(k) balance at the end of the calendar year containing t and the amount of all withdrawals made by the individual from t to the end of the calendar year containing t . We then reduce the individual's synthetic 401(k) balance by this fraction at the end of t .¹⁸

Employees can also take out loans against their 401(k) balance. Our analysis ignores 401(k) loans, which is equivalent to assuming that these loans are all repaid in full along with 5% annual interest. This biases our results toward finding a slightly larger treatment effect.¹⁹

4. Withdrawals at employment separation

Upon separating from a firm, the employee's money in the 401(k) sponsored by that firm becomes withdrawable at any time for any reason. Employees must pay ordinary income tax on any withdrawals that are not rolled over within 60 days to another 401(k) or an Individual Retirement Account (IRA). Employees younger than 59½ also usually pay a tax penalty equal to 10% of the withdrawal unless the money is rolled over. Individuals can request that the 401(k) administrator directly roll over the withdrawal into another retirement account, in which case the withdrawn amount never passes through the individual's bank account.

¹⁷ The match threshold increased only for some employees at firms H and I, but our data do not explicitly identify which employees received this increase. We can adjust the match threshold that applies to control employees one year before the actual match threshold change only for those whom we observe contributing at or above the new match threshold after the actual match threshold change. In the final year we observe an employee, we assume that she has her last known match structure.

¹⁸ Firm H is missing withdrawals data in 2009 and 2012. We assume that its employees made zero in-service withdrawals in those years.

¹⁹ Lu et al. (2017) report that 90% of 401(k) loans are fully repaid. To the extent that the counterfactual assumption of full repayment generates a bias, it probably makes our treatment effect estimates larger because employees with weak savings motives are probably more likely to default on 401(k) loans.

We observe in our data the amount within each calendar year that is directly rolled over and the amount that is paid out to the individual. We do not observe if an individual who has had a withdrawal paid out to them subsequently rolls over the withdrawal on their own. Such instances are likely to be rare because it is more convenient to execute a direct rollover, and direct rollovers are not subject to tax withholding, unlike withdrawals paid to an individual.

The threshold of \$1,000 of balances at separation is significant because during our sample period, if the employee's 401(k) balance at separation was less than \$1,000, the employer could—without the employee's consent—completely liquidate the employee's 401(k) account and send the proceeds to the employee via check. Although the individual can in principle roll over the cash distribution themselves, small-dollar withdrawals from retirement accounts are in practice especially unlikely to be rolled over (Poterba, Venti, and Wise, 1998; Argento, Bryant, and Sabelhaus, 2015). Therefore, whether one accumulates balances greater than \$1,000 before leaving the firm plays a large role in determining whether one's 401(k) dollars stay within the retirement savings system. Balances between \$1,000 and \$5,000 could also be unilaterally moved out of the 401(k) by the employer into an IRA of the employer's choice. Balances above \$5,000 had to be retained in the employer's 401(k) indefinitely until the employee chose to withdraw them.²⁰

Figure 2 shows how the post-separation cash leakage rate in the year of employment separation varies with our estimate of 401(k) balance at separation. We estimate this balance as the sum of post-separation cash withdrawal dollars during the calendar year of separation, post-separation rollover dollars during the calendar year of separation, and 401(k) balance on December 31 of the separation year. The cash leakage rate is post-separation cash withdrawal dollars during the calendar year of separation divided by estimated 401(k) balance at separation. The sample pools data on all treatment and control cohort employees at all firms (except for Firm H, which is missing withdrawals data in 2009 and 2012) who separate in a July, August, or September within our sample period. We restrict the separation months because compulsory cash distributions are typically enacted three to four months after an employee's separation date (see Appendix A). Individuals separating in later months have not been potentially subject to a compulsory cash-out before the closest year-end. For individuals separating in earlier months, we estimate balances at

²⁰ The \$5,000 threshold was increased to \$7,000 starting in 2024 by the SECURE 2.0 Act.

separation with greater error if they do not withdraw or roll over all their balances, since we only observe balances at each year-end.²¹

Consistent with firms' ability to compulsorily cash out balances under \$1,000, we see in Figure 2 that the cash leakage rate for such balances is almost 100%. The fact that the cash leakage rate for balances under \$1,000 is not exactly 100% is mostly due to employees proactively rolling over their balances to an IRA.²² Cash leakage rates drop discretely to around 50% right above the \$1,000 balance threshold, and then continue to decline gradually. Even for balances between \$10,000 and \$20,000, which are well above the threshold that allows individuals to keep their money in the original 401(k), the leakage rate is about 35%.

Because of the critical role of balance at separation in determining cash leakage, we reduce a separating employee's synthetic balances in the retirement savings system by a cash leakage percentage that depends on the employee's synthetic balance at separation. Appendix A presents evidence that any cash withdrawals from the 401(k) tend to be made soon after separation. Therefore, when an individual separates, we immediately reduce their synthetic balance by a cash leakage percentage and do not impose subsequent reductions. We treat rollover balances as remaining in the retirement savings system indefinitely, as our data contain no information on what happens to these balances after they leave the sponsoring employer's 401(k). This analytic choice likely biases our treatment effect estimates upward, since individuals induced to contribute more to the 401(k) due to an automatic savings policy are likely to have weak savings motives, and are thus more likely to withdraw rollover balances.

We compute the average leakage rate for July, August, and September leavers in the calendar year of separation within each of the balance bins shown in Table 4 separately for the treatment and control cohorts.²³ Conditional on balance size, leakage rates are similar between the cohorts and often statistically indistinguishable from each other. Where there are differences, the treatment cohort's leakage rates are somewhat higher, consistent with automatic policies bringing people with low savings motivation into the 401(k) participant pool. We reduce each separating employee's synthetic 401(k) balance by the leakage rate in the cell that matches their cohort and

²¹ If the individual withdraws or rolls over their entire balances, we know the amount of these distributions, so we know their balances at (or close to) separation. For individuals who keep balances in the 401(k), balances at year-end would differ from balances at separation because of investment gains and losses.

²² One-third of individuals with leakage rates less than zero retained balances in the 401(k).

²³ Online Appendix Table 3 shows that variation in balance is a far more important correlate of the leakage rate than variation in employee age.

synthetic balance at separation.²⁴ Because the synthetic balances described in equation (1) are normalized by salary, we multiply normalized balance at separation by the individual’s first observed salary (deflated to dollars at December 31 within the control cohort’s hire window) to obtain synthetic balances in absolute dollars. If salary data are missing for an employee’s entire tenure, then we use the median first-observed salary for employees hired in the two-year window centered on the automatic policy introduction date.

The penultimate row of Table 3 shows that at four firms, we need to impute salary for almost nobody when calculating the final synthetic 401(k) balance, but among the remaining five firms, we impute salary for between 3% and 17% of employees. The number of employees missing all salary observations is larger (shown in the last row), but we do not need to impute salaries for many of them because they never have positive contribution rates. The fact that these never-contributing employees are missing salaries does reduce our sample size later when we run regressions that control for salary.

5. Simulating future outcomes for employees who do not separate by 60 months of hire

For individuals who do not separate from the job by 60 months of tenure (i.e., the end of our observation period), we simulate their future separation date using the empirically observed firm-specific average monthly rate at which control and treatment employees (pooled together) leave the firm from tenure months 49 to 60.²⁵ We randomly assign these remaining employees to leave each month beyond tenure month 60 at this separation rate until every employee has separated from the firm. Using a similar approach, we simulate at firms with auto-escalation whether a remaining employee who is below the auto-escalation cap enters or leaves auto-escalation each year.²⁶ We impose a cash leakage rate upon simulated employment separation

²⁴ For this calculation, we use the synthetic balance at the moment of separation, without imposing a multi-month delay during which the account accrues capital gains before evaluation.

²⁵ We estimate this rate as the number of employees who depart the firm in that period divided by the total number of employee-months in which employees were working at the firm in that period.

²⁶ At firms where the difference between the initial default contribution rate and the auto-escalation cap is less than 5% of income, let T equal this difference (in integer units). At other firms, let $T = 5$. Define a “prompt” as a date on which one could potentially auto-escalate. We estimate separately for each firm, for employees below the auto-escalation cap, the probability of not auto-escalating on their T th prompt conditional on having auto-escalated on their $(T - 1)$ th prompt, and the probability of auto-escalating on their T th prompt conditional on not auto-escalating on their $(T - 1)$ th prompt. We use these entry and exit probabilities in our simulation. At Firm H, due to missing contribution rate data, we use the control cohort’s conditional probability of entry or exit at the second prompt and the treatment cohort’s conditional probability of entry or exit at the first prompt. Note that we do not directly observe whether

using the same procedure as previously described. Finally, to make simulation noise negligible, we run 100 simulations for each individual who does not separate from employment by 60 months of tenure, and we consider the average outcome across these 100 simulations to be the “observation” for this individual.

6. Conversion to equivalent constant contribution rate

The final outcome of interest for each individual i is the constant monthly contribution rate c_i^* during their tenure at the firm that would result in the same final post-leakage synthetic balance at employment separation (computed from actually observed data or partially simulated data), b_{iT} . Using the equations $b_{i1} = c_i^*/12$ and $b_{it} = 1.05^{1/12}b_{i,t-1} + c_i^*/12$, we get the expression for the “equivalent constant contribution rate”:

$$c_i^* = 12 \cdot \frac{1 - 1.05^{1/12}}{1 - 1.05^{T/12}} b_{iT} \quad (2)$$

C. Estimating equation

To estimate the effect of the automatic policy at each firm, we estimate the following regression separately by firm:

$$c_i^* = \beta_0 + \beta_1 T_i + \gamma W_i + u_i \quad (3)$$

where T_i is an indicator for whether individual i was hired after the automatic policy was implemented, and u_i is a residual term. In regressions that include additional controls, W_i is a vector containing a female indicator, the log of salary, a quadratic in age at hire, and a spline for maximum tenure achieved (right-censored at 60 months) with knot points every six months.²⁷ Robust standard errors are estimated.

employees are enrolled in auto-escalation; we infer this status from whether their contribution rate rises by one percentage point at the prompt.

²⁷ For simplicity, we right-censor the maximum-tenure-achieved control variable at 60 months in all our regressions, even when estimating treatment effects that are extrapolated from 12 months of tenure.

III. Results

A. Naïve extrapolation from first year of tenure, assuming 85% acceptance of auto-escalation default, 100% vesting, no withdrawals

Panel A of Table 5 shows estimates of treatment effects on equivalent constant contribution rates derived from comparisons of the treatment and control cohort averages with no additional control variables. The first number in each column comes from naively extrapolating from contribution rates observed up to one year of tenure, assuming no pre-retirement withdrawals and an 85% acceptance rate of auto-escalation in the treatment cohort at firms with auto-escalation. This approach is a natural way to extrapolate from the results of early pioneering papers in this literature (Madrian and Shea, 2001; Benartzi, Peleg, and Thaler, 2013).

In this first approach, the month 12 contribution rate of any individual still employed at the company at month 12 is assumed to remain in effect until month 60 unless the individual is at a firm with auto-escalation and increased their contribution rate by 1% on their first escalation date²⁸, in which case we assume that they will continue auto-escalating on schedule (with no other changes to their contribution rate) through month 60 or until their auto-escalation cap is hit. If fewer than 85% of treatment cohort individuals at a firm with auto-escalation are observed to take up auto-escalation on the first escalation date, we randomly assign additional treatment cohort individuals (preferentially choosing those below the auto-escalation cap²⁹) to accept auto-escalation so that the total take-up rate in that firm's treatment cohort is 85%. We assume that all matching contributions are 100% vested. Equivalent constant contribution rates are computed using contribution rates (actually observed until month 12 and extrapolated after month 12) either until employment separation for those who leave the firm before month 12 or until month 60 (for everyone else, regardless of their true separation date). To minimize the simulation error resulting from the random drawing of additional auto-escalation accepters, we repeat the above procedure

²⁸ We perform this extrapolation into future auto-escalation dates for both treatment and control cohort employees, since control employees could opt into auto-escalation. For employees whose first escalation date does not occur until their second year of tenure, we use whether they escalated on that first date—even though it occurs after tenure month 12—as the basis for our extrapolation.

²⁹ Only if there are not enough such employees to attain 85% take-up do we assign employees who are at or above the cap to auto-escalation. This latter group of employees does not end up auto-escalating unless they subsequently lower their contribution rate below the cap.

100 times, and use the average equivalent constant contribution rate for each individual across simulations as their single “observation.”

With this extrapolated measure, we estimate that at the four automatic enrollment firms (A through D), the automatic policy increases the equivalent constant contribution rate by 1.2%, 2.1%, 4.4%, and 0.3% of income per period, respectively. Figure 3 shows that at Firm A, the percent of active employees (i.e., those still employed at the firm) at the automatic enrollment default increases by 68 percentage points for the treatment cohort relative to the control cohort in tenure month 4 (the first month of 401(k) eligibility at this firm). This increase is almost exactly equal to the 72 percentage point increase in the fraction of employees with a positive contribution rate at month 4 (see Figure 4). However, the impact on 401(k) accumulation is muted because Firm A’s default contribution rate is a low 2%, so the average total employee plus match contribution rate (counting the entire match, whether or not it eventually vests fully) is only 1.6% of income higher in the treatment cohort than the control cohort among active employees at month 4 (see Figure 5). Furthermore, by month 4, about 20% of the cohort has already left the firm and thus is never subject to automatic enrollment (see Figure 6).

At Firm B, the increase in those contributing at the automatic enrollment default right after the opt-out deadline is smaller—only 30 percentage points, which again almost exactly matches the increase in the likelihood of a positive contribution rate. However, the default is 4% instead of 2%, causing the average total contribution rate among active employees to be 2.9% of income higher in the treatment cohort than the control cohort, and job attrition is much lower (4% by month 3), so the treatment effect is larger at Firm B than at Firm A.

The treatment effect is even larger at Firm C because it has a relatively high 4% default, there is a 64 percentage point difference between the fraction of the treatment versus control cohort contributing at the automatic enrollment default (as well as the fraction contributing at all) at month 1, and attrition is relatively low (2% at month 1 and 17% at month 12 for the treatment cohort). Although the automatic enrollment default is also the match threshold, few employees in the control cohort contribute at the match threshold, perhaps because the match is 0% vested before employees exceed three years of tenure.

Surprisingly, the automatic enrollment effect is the smallest at Firm D, which has the highest default contribution rate (5%). This is primarily because the control cohort has a high propensity to save. Choukhmane (2024) documents that companies that choose high default

contribution rates tend to have had high average contribution rates before automatic enrollment. The total contribution rate gap between the two Firm D cohorts among active employees is 2.3% of income at the opt-out deadline but falls rapidly to 0.6% of income by month 12. The treatment cohort continues to have significantly higher participation rates at higher tenures, so the similarity in average contribution rates suggests that automatic enrollment is dragging down the contribution rates of some employees who would have otherwise contributed a high amount.

Despite the small estimate at Firm D, on average across firms, automatic enrollment has a large estimated positive effect on contribution rates—2.0% of income—under this first approach to estimating the treatment effect.

We also estimate large effects from auto-escalation added on top of pre-existing automatic enrollment when extrapolating from the first 12 months of tenure and assuming an 85% take-up rate of auto-escalation. At Company E, the increase in the equivalent constant contribution rate is 2.4% of income, and at Company F, the increase is 1.7% of income. Figure 4 shows that the fraction of active employees with a positive contribution rate is high and nearly identical across treatment and control cohorts within each company, consistent with automatic enrollment being in place for both cohorts. In Figure 5, we see that the average total contribution rate across the first twelve months of tenure³¹ is, respectively, 0.8% and 0.1% of income higher in the treatment group than the control group at Companies E and F, reflecting in part the somewhat higher income in the treatment group than the control group at these firms. The fact that our extrapolated treatment effect is considerably larger than the total contribution rate gaps in the first twelve months indicates that most of the extrapolated treatment effect is coming from auto-escalation assumed to occur after the first year of tenure.

Simultaneously adding automatic enrollment and auto-escalation has even larger contribution effects when extrapolating from the first 12 months of tenure and assuming an 85% take-up rate of auto-escalation: 3.6%, 4.2%, and 2.6% of income at Companies G, H, and I, respectively. Relative to an opt-in enrollment regime with no default auto-escalation, the combination of these two policies dramatically raises participation rates (Figure 4) and generates sizable average total contribution rate increases in the first twelve months of tenure (2.3%, 3.0%, and 1.2% of income at G, H, and I, respectively, in Figure 5) that are smaller than the total

³¹ These averages equally weight each tenure month's average total contribution rate among active employees through tenure month 12.

extrapolated treatment effect because of the auto-escalation that is assumed to happen after the first year.

Panel B of Table 5 shows regression-adjusted treatment effect estimates, where the control variables are a female dummy, the log of salary, a quadratic in age at hire, and a spline for maximum tenure achieved (right-censored at 60 months) with knot points every six months. Focusing on the first number in each column, we see that most of the estimates are similar to those in Panel A. Firms B, D, and E have the largest changes. We have already noted that the treatment cohort at Firm E has significantly higher income than the control cohort, so it is not surprising that Firm E's regression-adjusted treatment effect (1.7%) is smaller than its unadjusted treatment effect (2.4%).³² Firm B also has a treatment cohort that is older and better-paid than its control cohort, causing its regression-adjusted treatment effect (1.6%) to shrink relative to its unadjusted effect (2.1%). Much of the increase in Firm D's estimate is due to employees with missing salary data being dropped from the regression. The median missing-salary employee leaves the firm after only one month of tenure, before automatic enrollment has fully kicked in. The unadjusted treatment effect using only employees with salary data is 1.0%, which is considerably larger than the 0.3% unadjusted treatment effect in the full sample and closer to the 1.5% regression-adjusted treatment effect among employees with salary data.

B. Extrapolation from first year of tenure, using actual acceptance of first auto-escalation, assuming 100% vesting, no withdrawals

The last five numbers in the second row of Table 5, Panel A are treatment effect estimates when we continue to extrapolate from the first year of tenure as in the first row, imposing 100% vesting of matching contributions and no withdrawals from the 401(k), but assuming an individual continues to auto-escalate only if they accept escalation on their first escalation date. The estimated effects at firms with auto-escalation shrink considerably relative to those in the first row. At Firms E and F, the effect of auto-escalation falls by 0.9% of income at both firms, a 37% and 51% relative decline, respectively. At the autoenrollment and escalation firms G, H, and I, the attenuation is

³² Firm E's treatment cohort is also older on average than its treatment cohort. However, the regression coefficients on age and the square of age indicate that age has a positive marginal effect on the equivalent constant contribution rate for only the older half of its sample.

smaller but significant—0.4%, 0.2%, and 0.5% of income, respectively, or a 10%, 6%, and 18% relative decline.

Why does the estimated treatment effect decline so dramatically at Firms E and F? The first two graphs in Figure 7 show the take-up rate of auto-escalation at these firms.³³ Because we do not directly observe an individual's auto-escalation election, we count an individual as having taken up auto-escalation if their contribution rate increases by 1% on the escalation date. This method will generate some false positives, but we can get a sense of the false positive rate by using the same method on the control cohort (setting their placebo escalation date sequence to start twelve months before the treatment group's actual escalation date sequence).³⁴ Only 30% of Firm E's treatment cohort and 27% of Firm F's treatment cohort that is still actively employed on the first escalation date accept the auto-escalation default. These percentages are much higher than for the control cohort (3% and 2%, respectively), but far below the 85% acceptance rates reported by Benartzi, Peleg, and Thaler (2013). The low take-up rates are not due to treatment cohort individuals actively choosing contribution rates that are above the auto-escalation cap; only 31% of treatment individuals below the cap just before the escalation date at Firm E escalate on the first escalation date, and this percentage is the same at Firm F.

The treatment effect declines are smaller at Firms G and H because a larger fraction of active treatment employees accept the first escalation—57% at Firm G and 67% at Firm H. Given that only 17% of employees at Firm I escalate at the first escalation date, one would expect that Firm I's estimated treatment effect would decline more than it does once we stop assuming an 85% take-up rate of auto-escalation. The reason the attenuation is small is that many of Firm I employees actively choose to contribute at or above the 6% auto-escalation cap (which happens to also be the match threshold), so assuming that these employees accept the auto-escalation default does not affect their contribution rate path. The take-up rate of auto-escalation among Firm I employees below the cap is 43%.

The regression-adjusted estimates in Table 5, Panel B are close to Panel A's unadjusted estimates, except for Firm E, whose adjusted estimate is 0.7% of income smaller than its raw estimate. We had already seen a smaller adjusted treatment effect for Firm E in the previous set of

³³ In Figure 7, if an individual's contribution rate is interpolated on an escalation date due to missing data, they are excluded from both the numerator and denominator of the acceptance rate calculation.

³⁴ Note that it is possible for control cohort employees to opt into automatic escalation, so it is not necessarily true that all control cohort members who are deemed to have experienced automatic escalation are false positives.

treatment effect estimates due to its treatment cohort having a higher average income than its control cohort.

C. Treatment effects including longer-horizon behavior, no withdrawals

In the third row of Table 5, Panel A, we present treatment effect estimates using actual contribution rates over the first 60 months of tenure, assuming all employer matches are 100% vested and assuming that no 401(k) withdrawals are made, either while employed or upon separation from the firm. Behavior beyond month 60 is simulated as described in Section II.B.5. Relative to the estimates that extrapolate from the first 12 months of tenure, the treatment effects shrink everywhere except at Firm B.

Figure 5 best summarizes why the treatment effects almost always shrink: contribution rate choices that employees make beyond one year of tenure. At two of the three automatic enrollment firms (A and C), the gap in the total contribution rate between treatment and control cohorts closes over time because the control cohort increases its contribution rate more quickly than the treatment cohort, a phenomenon that has also been documented by Choi et al. (2004) and Choukhmane (2024). Firm B is the exception where the contribution rates of both cohorts rise almost in parallel through tenure month 60, which is why its long-run treatment effect in the third row of Table 5, Panel A is slightly larger than in its extrapolated treatment effect in the second row.

In contrast, the long-run treatment effect at the five firms with auto-escalation shrinks not because the control cohort raises its contribution rate faster than the treatment cohort, but because the treatment cohort fails to pull away from the control cohort. The extrapolated treatment effect estimate assumed that the treatment cohort would continue to raise its contribution rate in tenure years 2-5 due to auto-escalation, while the control cohort's contribution rate would remain flat (similar to the assumptions of Benartzi and Thaler (2013)). Instead, we see that the treatment cohort's contribution rate does not rise much more quickly than the control cohort's, despite the treatment cohort being subject to default auto-escalation. Figure 7 shows that the treatment cohort's failure to pull away may be partly attributable to the acceptance of auto-escalation decaying at each successive escalation date. Consistent with this, Figure 3 illustrates that the fraction of active employees in the treatment cohort whose contribution rate equals that of a perfectly passive employee declines with tenure at each of the companies with auto-escalation. (At Firms F, G, and I, this fraction artificially increases at higher tenures when the passive contribution

rate rises to equal the match threshold, which had previously been actively chosen by many employees.) Nonetheless, a substantial fraction of the treatment cohort continues to accept escalation and remain on the passive contribution rate path. The similarity in average contribution growth rates between the treatment and control cohorts suggests that most of the treatment cohort individuals who accept auto-escalation would have raised their contribution at a similar rate in the absence of auto-escalation.

The fact that control cohort employees catch up to employees subject to automatic enrollment and keep up with employees subject to auto-escalation highlights the fact that 401(k) participants are less passive than commonly believed. The models of default effects presented by Choi et al. (2003) and Carroll et al. (2009) assume that individuals change their contribution rate at most once in a given company. In contrast, Figure 8 shows that in both the control and treatment cohorts at automatic enrollment firms, about half of individuals still employed at the firm at five years of tenure have switched their contribution rate at least twice *after* their fourth month of tenure (when the initial opt-out deadline has passed at all the firms), and about a third have switched their contribution rate at least three times.³⁵

The fourth row of Table 5, Panel A uses the actual vesting percentage of the employer match achieved by each employee, rather than assuming 100% vesting. Other than at Firms A, B, and E, which feature 100% immediate vesting, the treatment effects shrink further, reflecting the fact that many employees depart the firm before their match fully vests.

Although the magnitudes of the regression-adjusted treatment effects in the third and fourth rows of Table 5, Panel B differ a bit from the raw estimates, the findings are qualitatively the same. Going from the initial naïve estimates extrapolated from 12 months of contribution rates and assuming 100% vesting and 85% auto-escalation take-up at firms with auto-escalation to the estimates using 60 months of contribution rates and actual vesting, the raw treatment effect estimates fall by an average of 0.5% of income (26% decline relative to the average raw treatment effect) at the automatic enrollment firms, 1.0% of income (49% relative decline) at the auto-escalation firms, and 1.6% of income (46% relative decline) at the autoenrollment and escalation

³⁵ We focus on the automatic enrollment firms here as auto-escalation increases contribution rates over time for passive employees, complicating interpretation of how prone employees at those firms are to make active contribution rate switches.

firms. The respective regression-adjusted declines are 0.7% of income (33% relative decline), 1.2% of income (67% relative decline), and 2.0% of income (51% relative decline).

D. Comprehensive treatment effects

The fifth row of Table 5, Panel A contains our most comprehensive raw treatment effect estimates, now including the impact of withdrawals. The treatment effect falls even further at every firm, and becomes negative (but statistically indistinguishable from zero) at Firm D. The average treatment effect is 0.7% of income among the autoenrollment firms, 0.7% of income among the auto-escalation firms, and 0.8% of income among the autoenrollment and auto-escalation firms. These are 55%, 30%, and 57% declines, respectively, relative to the average raw treatment effect in the fourth row, which incorporates everything except withdrawals. They are 67%, 65%, and 77% declines, respectively, relative to the most naive average raw treatment effects at the top of each column.

Withdrawals while employed at our sample firms are negligible, so the changes in the treatment effect estimates are due almost entirely to withdrawals upon separation. Table 6 shows the average cash leakage rate that is applied to members of each cohort at each firm at separation. Averaging across cohorts and firms equally, the leakage rate is 42%. If every individual's accumulation were shrunk by 42% at separation, the net-of-leakage treatment effect would mechanically shrink relative to the most comprehensive no-leakage treatment effect by 42%. In reality, the average leakage rate of the treatment cohort exceeds the control cohort at every firm but one, causing the treatment effect to attenuate further. We previously saw in Table 4 that conditional on balance at separation, leakage rates do not differ greatly between the treatment and control cohorts. Therefore, the fact that the leakage rate is higher for the treatment cohort indicates that treatment individuals are more likely than the control cohort to accumulate small positive balances, which then leak at a higher rate. Table 6 shows that on average, equally weighting each firm, the treatment cohort withdraws 8.5 percentage points more of its synthetic balance than the control cohort upon separation.

The fifth row of Table 5, Panel B shows regression-adjusted comprehensive treatment effect estimates. The declines relative to the most naïve regression-adjusted treatment effects are also large: 72%, 87%, and 79% declines on average for the automatic enrollment, auto-escalation, and autoenrollment and escalation firms, respectively. The average regression-adjusted treatment

effect is 0.6%, 0.2%, and 0.8% of income at each respective group. These are effects of modest magnitude, equal to 19%, 4%, and 22% of the corresponding control cohort's average equivalent constant contribution rate, shown in Panel C.

E. Auto-escalation treatment effects with Vanguard acceptance rates

Clark and Young (2021) report that in the Vanguard universe, the acceptance rate of an auto-escalation default at the first, second, and third escalation date is 63%, 63%, and 60%, respectively. These numbers are significantly higher than the acceptance rate in our sample. In this subsection, we examine how our treatment effect estimates might change if our sample's auto-escalation acceptance rates matched those in the Vanguard universe.

If the auto-escalation acceptance rate of a firm's treatment cohort on its first escalation date is below the Vanguard first-escalation acceptance rate, we randomly select additional treatment cohort plan participants below the auto-escalation cap to accept auto-escalation until the treatment cohort's acceptance rate matches the Vanguard acceptance rate.³⁶ Participants who have been randomly assigned to accept auto-escalation are artificially escalated on each subsequent escalation date (so that someone who has been subject to n artificial escalations up to time t will have their time t contribution rate increased by n percentage points relative to their actual contribution rate at t) until they either hit the escalation cap or leave the company. If the auto-escalation acceptance rate of the treatment cohort on the second or third escalation date falls below the Vanguard acceptance rate on the same escalation date, we assign additional treatment cohort participants below the cap to accept auto-escalation on that date until the treatment cohort's acceptance rate matches the corresponding Vanguard rate. To reduce simulation noise, we repeat the above exercise 100 times, and we consider each individual's average outcome across these 100 simulations to be the "observation" for the individual.

Note that this exercise probably overstates the contributions that higher auto-escalation take-up would create, since it assumes that accepting auto-escalation on an escalation date

³⁶ If there are not enough participants below the cap to reach the Vanguard acceptance rate, we additionally choose participants at or above the cap. The actual acceptance rate p is calculated over all treatment cohort members still at the firm with sufficient non-missing contribution rate data around the escalation date to assess whether they escalated. Those who are at or above the escalation cap are counted as not accepting auto-escalation. Let N be the total number of treatment cohort members still at the firm on the escalation date, and P be the Vanguard acceptance rate. Then we randomly select $(P - p)N$ treatment cohort members to enter auto-escalation. We handle missing escalation acceptance data at Firm H by assuming that on the escalation date with the missing data, the individual's acceptance status is the same as it was on the prior escalation date.

increases one's contribution rate by a full 1% of income relative to the counterfactual of not accepting. We have previously seen that most of those who actually take up auto-escalation in our sample would have increased their contribution rates at a similar pace without auto-escalation. It is likely that in reality, additional take-up of auto-escalation would be disproportionately driven by individuals who are prone to increase their contribution rates anyway (Carroll et al., 2009).

The final rows of Table 5 Panels A and B show the new treatment effects at the auto-escalation and autoenrollment and escalation firms. Without regression adjustments, the average treatment effect is 1.3% of income at the auto-escalation firms and 1.1% of income at the autoenrollment and escalation firms. With regression adjustments, the respective averages are 0.9% and 1.1%. Unsurprisingly, the treatment effects rise when we assume higher auto-escalation take-up—more so at the auto-escalation firms than at the autoenrollment and escalation firms because actual escalation take-up is lower in the former group. Nevertheless, they are significantly lower than the most-naïve estimates that ignore endogenous contribution rate catch-up, job turnover, incomplete vesting, and withdrawals.

F. Treatment effects on distribution of contribution rates

Prior literature has found that automatic enrollment has a larger positive effect on the left tail of the savings distribution (Madrian and Shea, 2001; Choi et al., 2002, 2004; Beshears et al., 2008; Choukhmane, 2024). Online Appendix Figure 1 reports distributions of equivalent constant contribution rates in our setting. The graphs in the left column present equivalent constant contribution rates calculated using observed contributions rates over the first 60 months of tenure, assuming 100% vesting and no leakage (corresponding to the inputs into the third row of Table 5). Automatic enrollment strongly reduces the share of savers at a 0% equivalent constant contribution rate—25 percentage points in the average firm, with smaller shifts at Firms B and D, which had relatively higher control cohort participation rates. Auto-escalation as an add-on to pre-existing automatic enrollment shifts the mass of savings to the right for right-tail savers; the fraction with an equivalent constant contribution rate above what would be achieved by always contributing at the firm's initial default is 15 percentage points higher on average.³⁸ Firms that simultaneously introduce both autoenrollment and auto-escalation exhibit both patterns—a 41 percentage point

³⁸ In this comparison, we do not adjust for the fact that the treated cohort at the auto-escalation firms have higher incomes, and thus higher baseline propensities to save, on average.

average decrease at 0% and a 16 percentage point average increase at levels above the initial automatic enrollment default.³⁹

The graphs in the right column of Online Appendix Figure 1 show distributions of *comprehensive* equivalent constant contribution rates (corresponding to the inputs into the fifth row of Table 5), which account for incomplete vesting and leakage. We have shown that these additional factors attenuate the average treatment effect substantially. We see in the figure that this attenuation of treatment effects occurs across the distribution. For automatic enrollment firms on average, incorporating leakage shrinks to 12 percentage points the amount by which automatic enrollment reduces the frequency of 0% equivalent constant contribution rates. For auto-escalation firms, incorporating leakage and non-vesting reduces the rightward shift in the right tail of savings rates; the increase in equivalent contribution rates exceeding the initial default level is only 8 percentage points. Firms that simultaneously introduce automatic enrollment and auto-escalation show both patterns of attenuation once non-vesting and leakage are accounted for—only a 14 percentage point decrease at 0% and a 7 percentage point increase above the initial automatic enrollment default.

IV. Conclusion

Automatic savings policies have been widely adopted in part because of the empirical evidence that they increase retirement savings accumulation. When we jointly incorporate several medium- and long-run dynamic factors, we find that the magnitude of the estimated savings effect is substantially diminished. Nevertheless, we document that automatic savings policies have a positive statistically significant impact on measured savings. The effects we estimate are modest in magnitude because of the increase in savings over time by those who are not subject to automatic policies, vesting requirements with respect to employer matching contributions, high employee turnover rates (which undermines auto-escalation and vesting), the low acceptance of auto-escalation defaults, and the high rate of cash distributions (i.e., leakage) upon job separation. Automatic savings policies are likely cost-effective from an impact-to-cost ratio perspective (Benartzi et al., 2017). However, the current paper does not attempt to provide a (structural)

³⁹ We use as the cutoff for “above the initial default” 0.5 percentage points above the equivalent contribution rate associated with always contributing at the initial default under the most common match structure in the most recent period. This translates to 12.5% at Firm E, 6.5% at Firm F, 4.5% at Firm G, and 6.5% at Firms H and I.

welfare analysis of automatic savings policies. If policymakers in the U.S. wish to generate larger changes in savings rates, coupling nudges with reduced liquidity of retirement savings before retirement and/or increasing compulsory savings may be more effective at raising savings and household welfare (Chetty et al., 2014; Beshears et al., 2024; Chater and Loewenstein, 2023). In other countries that use defined contribution savings systems, illiquidity and compulsory savings are the norm (e.g., see Beshears et al., 2015).

Appendix A. Timing of withdrawals after employment separation

Appendix Figure 1 shows that the cumulative cash leakage rate swiftly approaches its asymptote after job separation among treatment and control cohort employees in all firms except Firm H who separate in a July, August, or September.

We construct the cumulative cash leakage rate variable as follows in order to bound its value between 0% (if no cash withdrawals are taken) and 100% (if the entire 401(k) balance is taken as a cash withdrawal) despite the fact that withdrawals can happen at various dates, there are capital gains and losses between those dates, and we only observe 401(k) balances at each year-end.

Define the cash leakage rate of individual i on day $d \in \{0, 1, 2, \dots, 365\}$ of year y as

$$l_{idy} = \frac{w_{idy}}{B_{iy} + R_{iy} + W_{iy}} \quad (\text{A1})$$

where w_{idy} is the dollars disbursed as a cash withdrawal to i on day d of year y , B_{iy} is i 's (actual, not synthetic) 401(k) balance at the end of year y , R_{iy} is the total dollars withdrawn by i as rollovers in year y , and W_{iy} is the total dollars of cash withdrawals by i in year y . The denominator $B_{iy} + R_{iy} + W_{iy}$ is what i 's 401(k) balance would be at the end of y if there were no withdrawals, rollovers, or capital gains during y . Similarly, define the rollover rate as

$$r_{idy} = \frac{\rho_{idy}}{B_{iy} + R_{iy} + W_{iy}} \quad (\text{A2})$$

where ρ_{idy} is the dollars disbursed as a rollover to i on day d of year y .

In the calendar year of separation, i 's cumulative cash leakage rate from his day of separation s through day $d \geq s$ is

$$\mathcal{L}_{idy} = \sum_{t=s}^d l_{ity} \quad (\text{A3})$$

and the cumulative rollover rate is

$$\mathcal{R}_{idy} = \sum_{t=s}^d r_{ity} \quad (\text{A4})$$

In subsequent years, the cumulative cash leakage rate is defined as a weighted sum of the cumulative cash leakage rate at the end of the prior calendar year and the sum of the daily cash leakage rate in the current calendar year:

$$\mathcal{L}_{idy} = \mathcal{L}_{i,365,y-1} + (1 - \mathcal{L}_{i,365,y-1} - \mathcal{R}_{i,365,y-1}) \sum_{t=1}^d l_{ity} \quad (\text{A5})$$

The value of \mathcal{L}_{idy} is bounded between 0% and 100%. The cumulative rollover rate is similarly defined as

$$\mathcal{R}_{idy} = \mathcal{R}_{i,365,y-1} + (1 - \mathcal{L}_{i,365,y-1} - \mathcal{R}_{i,365,y-1}) \sum_{t=1}^d r_{ity} \quad (\text{A6})$$

If there are never any capital gains or losses, then \mathcal{L}_{idy} equals the sum of cash distribution dollars from separation to day d of year y divided by 401(k) balances at separation. If there are capital gains or losses and withdrawals always happen at calendar year-ends, then $\mathcal{L}_{i,365,y}$ is the amount that withdrawn cash would have been worth if left to appreciate inside the 401(k) until the end of y divided by what the 401(k) balance would have been at the end of y in the absence of withdrawals and rollovers.

To see why the above claims are true, let year 0 be the calendar year of job separation. Cumulative cash leakage in year 1 is

$$\mathcal{L}_{id1} = \mathcal{L}_{i,365,0} + (1 - \mathcal{L}_{i,365,0} - \mathcal{R}_{i,365,0}) \sum_{t=1}^d l_{it1} \quad (\text{A7})$$

$$= \frac{W_{i0}}{B_{i0} + R_{i0} + W_{i0}} + \frac{B_{i0}}{B_{i0} + R_{i0} + W_{i0}} \sum_{t=1}^d \frac{w_{id1}}{B_{i1} + R_{i1} + W_{i1}} \quad (\text{A8})$$

If there are no capital gains, then $B_{i0} = B_{i1} + R_{i1} + W_{i1}$, so (A8) is equal to

$$\frac{W_{i0} + \sum_{t=1}^d w_{it1}}{B_{i0} + R_{i0} + W_{i0}} \quad (\text{A9})$$

which is the sum of all cash withdrawals from separation to day d of year 1 divided by 401(k) balance at separation.

If 401(k) balances at the end of year 0 have a gross return of g_1 during year 1 and transactions during years 0 and 1 only occur at the very end of the year, then $B_{i0}g_1 = B_{i1} + R_{i1} + W_{i1}$. In this case, the expression from (A8) for $\mathcal{L}_{i,365,1}$ is equal to

$$\frac{W_{i0}g_1 + W_{i1}}{(B_{i0} + R_{i0} + W_{i0})g_1} \quad (\text{A10})$$

which is what cash withdrawals from separation to the end of year 1 would be worth at the end of year 1 if they had been left in the 401(k) divided by what the 401(k) would have been worth at the end of year 1 in the absence of cash withdrawals and rollovers.

The expressions for cumulative cash leakage in subsequent years have the same interpretations. In the absence of capital gains, the expression for $y > 1$ is

$$\frac{\sum_{v=0}^{y-1} W_{iv} + \sum_{t=1}^d w_{ity}}{B_{i0} + R_{i0} + W_{i0}} \quad (\text{A11})$$

and, with capital gains and withdrawals only at year-ends, is

$$\frac{\sum_{v=0}^{y-1} W_{iv} (\prod_{q=v+1}^y g_q) + W_y}{(B_{i0} + R_{i0} + W_{i0}) \prod_{v=1}^y g_v} \quad (\text{A12})$$

Appendix B. Acceptance of default auto-escalation in a larger sample of firms

Using plan documents and a survey Alight conducted of its clients, we identify 20 retirement savings plans in our data that implemented default auto-escalation, have sufficiently high data quality, and for which our data allow us to observe the first three escalation dates.⁴⁰ We additionally include Firm H despite its only having data on the first two escalation dates because it is in our main text's analysis. We restrict our sample to employees hired in the 365 days after the policy's introduction who were not rehires. Some firms have multiple retirement savings plans and do not implement default auto-escalation in all of them; in these instances, we exclude employees who are ineligible for the plan with default auto-escalation. Many firms require an employee to have a minimum amount of tenure on an escalation date in order to be auto-escalated. We consider the first escalation date on which the employee is eligible for auto-escalation as their "first" escalation date, even if other employees at the firm were escalated on an earlier calendar date.

⁴⁰ Firm C is included in this sample because it implemented default auto-escalation in 2008. In the main text, we study Firm C as an autoenrollment firm because it implemented automatic enrollment in 2003.

We count an employee as having accepted auto-escalation on an escalation date if their contribution rate at the end of the escalation month is 1 percentage point higher (the auto-escalation increase amount at all of our sample firms) than at the end of the prior month. An important caveat to this analysis is that we are missing contribution rate data for some employees in the escalation month and the prior month. We compute escalation acceptance rates excluding these employees from both the numerator and denominator. We also exclude from the numerator and denominator employees who have separated from the firm prior to the escalation date and employees whose escalation decision we do not observe on that date because it occurs beyond the sample period covered by our data.

Appendix Table 1 shows, for each firm, the escalation acceptance rate at each escalation date, as well as the percent of employees still at the firm in those months whose contribution rates are missing. On average, the acceptance rate is 43% on the first escalation date, 36% on the second date, and 29% on the third date. Conditional on being below the auto-escalation cap in the month before the escalation date, the average acceptance rate is 54% on the first date, 48% on the second date, and 39% on the third date.

References

- Alight, 2019. “What do workers do with their retirement savings after they leave their employers? A deep dive into post-termination behavior, 2008-2017.” Alight Solutions white paper.
- Altmann, Steffen, Andreas Gruenwald, and Jonas Radbruch, 2024. “The double dividend of attention-releasing policies.” CESifo Working Paper No. 11069.
- Argento, Robert, Victoria L. Bryant, and John Sabelhaus, 2015. “Early withdrawals from retirement accounts during the Great Recession.” *Contemporary Economic Policy* 33, 1-16.
- Benartzi, Shlomo, John Beshears, Katherine L. Milkman, Cass R. Sunstein, Richard H. Thaler, Maya Shankar, Will Tucker-Ray, William J. Congdon, and Steven Galing, 2017. “Should governments invest more in nudging?” *Psychological Science* 28, 1-15.
- Benartzi, Shlomo, Ehud Peleg, and Richard H. Thaler, 2013. “Choice architecture and retirement savings plans.” In Eldar Shafir, ed., *The Behavioral Foundations of Public Policy*. Princeton, NJ: Princeton University Press, 245-263.
- Benartzi, Shlomo, and Richard H. Thaler, 2013. “Behavioral economics and the retirement savings crisis.” *Science* 339(6124), 1152-1153.
- Beshears, John, Matthew Blakstad, James J. Choi, Chris Firth, John Gathergood, David Laibson, Richard Notley, Jesal D. Sheth, Will Sandbrook, and Neil Stewart, 2024. “Does pension automatic enrollment increase debt? Evidence from a large-scale natural experiment.” NBER Working Paper 32100.

- Beshears, John, James J. Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte C. Madrian, 2024. "Optimal illiquidity." *Journal of Financial Economics*, forthcoming.
- Beshears, John, James J. Choi, Joshua Hurwitz, David Laibson, and Brigitte C. Madrian, 2015. "Liquidity in Retirement Savings Systems: An International Comparison." *American Economic Review Papers and Proceedings* 105(5), pp. 420-425.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, 2008. "The importance of default options for retirement saving outcomes: Evidence from the United States." In Stephen J. Kay and Tapen Sinha, eds., *Lessons from Pension Reform in the Americas*. Oxford: Oxford University Press, pp. 59-87.
- Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and William L. Skimmyhorn, 2022. "Borrowing to save? The impact of automatic enrollment on debt." *Journal of Finance* 77, pp. 403-447.
- Beshears, John, James J. Choi, David Laibson, and Peter Maxted, 2022. "Present bias causes and then dissipates auto-enrollment savings effects." *AEA Papers and Proceedings* 112, pp. 136-141.
- Blumenstock, Joshua, Michael Callen, and Tarek Ghani, 2018. "Why do defaults affect behavior? Experimental evidence from Afghanistan." *American Economic Review* 108, 2868-2901.
- Brown, Alexander L., Daniel Grodzicki, and Paolina C. Medina, 2023. "When nudges spill over: Student loan use under the CARD Act." Mimeo.
- Carranza, Guillermo, and Aaron Goodman, 2024. "Retention or regressivity? The empirical effects of 401(k) vesting schedules." Vanguard working paper.
- Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2009. "Optimal defaults and active decisions." *Quarterly Journal of Economics* 124, 1639-1674.
- Chater, Nick, and George Loewenstein, 2023. "The i-frame and the s-frame: How focusing on individual-level solutions has led behavioral public policy astray." *Behavioral and Brain Sciences* 46, e147.
- Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen, 2014. "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark." *Quarterly Journal of Economics* 129, 1141-1219.
- Choi, James J., 2015. "Contributions to defined contribution pension plans." *Annual Review of Financial Economics* 7, 161-178.
- Choi, James M., David Laibson, Brigitte C. Madrian and Andrew Metrick, 2002. "Defined contribution pensions: Plan rules, participant decisions, and the path of least resistance." In James Poterba, ed., *Tax Policy and the Economy* 16, pp. 67-114.
- Choi, James M., David Laibson, Brigitte C. Madrian and Andrew Metrick, 2003. "Optimal defaults." *American Economic Review Papers and Proceedings* 93, 180-185.
- Choi, James J., David Laibson, Brigitte C. Madrian and Andrew Metrick, 2004. "For better or for worse: Default effects and 401(k) savings behavior." In David A. Wise, ed., *Perspectives on the Economics of Aging*. Chicago: University of Chicago Press, pp. 81-121.

- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2009. “Reinforcement learning and savings behavior.” *Journal of Finance* 64, 2515-2434.
- Choukhmane, Taha, 2024. “Default options and retirement savings dynamics.” *American Economic Review*, forthcoming.
- Clark, Jeffrey W., and Jean A. Young, 2021. “Automatic enrollment: The power of the default.” Vanguard Research. https://institutional.vanguard.com/iam/pdf/ISGAE_022020.pdf
- DellaVigna, Stefano, and Elizabeth Linos, 2022. “RCTs to scale: Comprehensive evidence from two nudge units.” *Econometrica* 90, 81-116.
- Department for Work & Pensions, 2022. “Workplace pension participation and savings trends of eligible employees: 2009 to 2021.” <https://www.gov.uk/government/statistics/workplace-pension-participation-and-savings-trends-2009-to-2021/workplace-pension-participation-and-savings-trends-of-eligible-employees-2009-to-2021>
- Derby, Elena, Kathleen Mackie, and Jacob Mortenson, 2023. “Worker and spousal responses to automatic enrollment.” *Journal of Public Economics* 223, 104910.
- Guttman-Kenney, Benedict, Paul Adams, Stefan Hunt, David Laibson, Neil Stewart, and Jesse Leary, 2024. “The semblance of success in nudging consumers to pay down credit card debt.” *American Economic Journal: Economic Policy*, forthcoming.
- Jachimowicz, Jon M., Shannon Duncan, Elke U. Weber, and Eric J. Johnson, 2019. “When and why defaults influence decisions: a meta-analysis of default effects.” *Behavioural Public Policy* 3, 159-186.
- Lu, Timothy (Jun), Olivia S. Mitchell, Stephen P. Utkus, and Jean A. Young, 2017. “Borrowing from the future? 401(k) plan loans and loan defaults.” *National Tax Journal* 70, 77-110.
- Madrian, Brigitte C., and Dennis F. Shea, 2001. “The power of suggestion: Inertia in 401(k) participation and savings behavior.” *Quarterly Journal of Economics* 116, pp. 1149-1187.
- Medina, Paolina C., 2021. “Side effects of nudging: Evidence from a randomized intervention in the credit card market.” *Review of Financial Studies* 34, 2580-2607.
- Medina, Paolina C., and Michaela Pagel, 2023. “Does saving cause borrowing? Implications for the co-holding puzzle.” Mimeo.
- OECD, 2021. *Pensions at a Glance 2021: OECD and G20 Indicators*. Paris: OECD Publishing. <https://doi.org/10.1787/ca401ebd-en>.
- Poterba, James J., Steven F. Venti, and David A. Wise, 1998. “Lump-sum distributions from retirement savings plans: Receipt and utilizations.” In David A. Wise, editor, *Inquiries in the Economics of Aging*, 85-108.
- Thaler, Richard H., and Shlomo Benartzi, 2004. “Save More Tomorrow™: Using behavioral economics to increase employee saving.” *Journal of Political Economy* 112, S164-S187.
- Vanguard, 2013. *How America Saves 2013: A Report on Vanguard 2012 Defined Contribution Plan Data*. <https://info.rch1.com/hubfs/Publications/Vanguard/2013%20Vanguard%20study%20How%20America%20Saves%202013.pdf>

Vanguard, 2023. *How America Saves 2023*.

<https://institutional.vanguard.com/content/dam/inst/iig-transformation/has/2023/pdf/has-insights/how-america-saves-report-2023.pdf>

Zook, David, 2023. “How do retirement plans for private industry and state and local government workers compare?” *Beyond the Numbers: Pay & Benefits* 12(1) (U.S. Bureau of Labor Statistics).

Zhong, Mingli, 2021. “Optimal default retirement savings policies: Theory and evidence from OregonSaves.” Pension Research Council Working Paper 2020-1.

Table 1. Summary of automatic policy treatment effects

This table shows estimates of automatic savings policy treatment effects, averaged across firms for a given policy, on equivalent constant contribution rates using different methodologies. Each successive row adds one methodological refinement relative to the prior row. All estimates control for demographics and maximum tenure achieved, as described in Table 5. Standard errors are shown in parentheses, calculated assuming each company's estimated treatment effect is independent of the other companies' estimated treatment effects. That is, let $\sigma_1, \dots, \sigma_n$ be the standard errors of the n companies' treatment effects used to compute the average treatment effect. We compute the standard error of this average as $(\sqrt{\sigma_1^2 + \dots + \sigma_n^2})/n$.

	Autoenrollment	Auto-escalation	Both autoenrollment and escalation
Extrapolate from 1 year of tenure, assume 85% auto-escalation take-up	--	1.82% (0.09)	3.50% (0.11)
Extrapolate from 1 year of tenure	2.20% (0.23)	0.82% (0.09)	3.12% (0.11)
Extrapolate from 5 years of tenure	1.84% (0.22)	0.66% (0.09)	2.41% (0.11)
Incorporate incomplete vesting of employer contributions	1.48% (0.20)	0.60% (0.08)	1.84% (0.10)
Incorporate withdrawals (comprehensive estimate)	0.62% (0.16)	0.23% (0.07)	0.81% (0.08)
Use Vanguard's auto-escalation take-up rate	--	0.87% (0.07)	1.11% (0.08)

Table 2. 401(k) plan features

Firm	Policy intro date	Default contribution rate	Employer match and nonelective contributions	Employer contribution vesting schedule
Panel A: Autoenrollment firms				
A	Jul 1, 2005	2%	100% of first 4% of income contributed	100% immediately
B	Apr 1, 2008	4%	100% of first 4% of income contributed	100% immediately
C	Jan 1, 2003	4%	100% of first 4% of income contributed	0% before tenure year 4, 100% at tenure year 4
D	Jun 2, 2006	5%	75% of first 2% of income contributed, 50% of next 3% of income contributed	0% before tenure year 4, 100% at tenure year 4
Panel B: Auto-escalation firms				
E	Jan 1, 2011	6% initial, escalate up to 15%	100% of first 6% of income contributed	100% immediately
F	Jan 1, 2011	3% initial, escalate up to 10%	100% of first 6% of income contributed for most employees. 50% or 75% match rate for ~15% of employees in 2010-2011, ~5% in 2012-2017	0% before tenure year 2, 40% at tenure year 2, 60% at tenure year 3, 80% at tenure year 4, 100% at tenure year 5
Panel C: Autoenrollment and escalation firms				
G	Jan 1, 2010	2% initial, escalate up to 6%	100% of first 6% of income contributed	0% before tenure year 4, 100% at tenure year 4
H	Jan 1, 2006	3% initial, escalate up to 15%	<i>Before Apr 1, 2008:</i> 50% of first 3% of income contributed <i>Starting Apr 1, 2008:</i> 100% of first 3% of income contributed*	0% before tenure year 4, 100% at tenure year 4
I	Jul 1, 2007	3% initial, escalate up to 6%	<i>Before Apr 1, 2009:</i> 80% of first 6% of income contributed if tenure < 5 years. 100% of first 6% of income contributed if tenure ≥ 5 years <i>Starting Apr 1, 2009:</i> 100% of first 6% of income contributed, plus 1% of income nonelective contribution. About 5% of employees have 7% match threshold.	0% immediately, 20% at tenure year 1, 40% at tenure year 2, 60% at tenure year 3, 80% at tenure year 4, 100% at tenure year 5

* A small number of Firm H employees had a match threshold between 4% and 9% during the sample period. Most significantly, a match threshold of 4% applied to about 5% of employees in 2010 and about 10% in 2011.

Table 3. Hire cohort characteristics

The top two sections show average characteristics, employment attrition rates, and employee counts in the control and treatment cohorts. Salary (deflated using CPI-U to the December 31 within the control cohort's hire window) is measured at hire when available; otherwise, it is the employee's first observed salary if one is available, or the median first-observed salary of everybody in the firm hired within 365 days before or after the policy implementation date. The penultimate section shows p -values from tests of equality across the two cohorts. The final section shows the fraction of employee-months (counting only months in which the individual is employed by the sample firm) where contribution rates are imputed, the fraction of employees whose salary is imputed to be the firm-wide median salary for the purposes of computing their final synthetic 401(k) balance, and the fraction of employees whose salary we never observe.

	Autoenrollment firms				Auto-escalation firms		Autoenrollment and escalation firms		
	Firm A	Firm B	Firm C	Firm D	Firm E	Firm F	Firm G	Firm H	Firm I
<i>Control cohort</i>									
Age at hire	29.8	32.3	33.2	34.8	35.9	37.2	33.0	35.2	33.6
Female	65.5%	49.3%	24.0%	33.1%	19.2%	23.5%	60.7%	77.6%	40.2%
Salary	\$23,748	\$55,333	\$56,606	\$49,826	\$40,930	\$59,671	\$29,831	\$31,471	\$57,305
5-year attrition	82.0%	57.5%	59.8%	82.9%	53.7%	67.7%	78.0%	71.2%	45.2%
Employees	13,275	793	445	245	1,963	5,137	1,697	30,912	1,470
<i>Treatment cohort</i>									
Age at hire	29.7	33.9	34.6	36.1	37.8	37.2	32.0	34.9	33.9
Female	66.8%	42.8%	21.1%	28.7%	22.5%	23.0%	58.9%	78.4%	38.8%
Salary	\$23,396	\$60,357	\$52,853	\$39,755	\$46,701	\$64,715	\$31,528	\$34,655	\$58,230
5-year attrition	80.1%	50.9%	55.5%	90.0%	58.3%	62.4%	79.9%	70.0%	43.3%
Employees	13,691	442	398	251	5,510	7,029	2,388	31,565	1,156
<i>Tests of equality of means or proportions between cohorts</i>									
Age at hire	0.599	0.011	0.049	0.159	< 0.001	0.786	0.005	0.003	0.372
Female	0.018	0.027	0.308	0.293	0.002	0.483	0.254	0.014	0.451
Salary	0.160	0.035	0.291	< 0.001	< 0.001	< 0.001	0.298	< 0.001	0.477
5-year attrition	< 0.001	0.026	0.213	0.020	< 0.001	< 0.001	0.127	0.002	0.348
<i>Imputed or missing observations</i>									
Contrib. imputed	9.3%	0.4%	5.3%	2.3%	2.7%	2.4%	1.8%	25.2%	0.8%
Salary imputed	0.0%	0.0%	2.8%	17.1%	4.8%	6.5%	0.0%	12.6%	0.5%
Salary never observed	2.5%	0.4%	14.6%	32.1%	6.8%	15.2%	2.3%	28.6%	2.3%

Table 4. Cash leakage rates by 401(k) balance at employment separation

This figure shows the cash leakage rate in the year of employment separation by our estimate of 401(k) balance at separation. The sample is treatment or control cohort employees who separated in July, August, or September at all firms except Firm H, which is missing withdrawals data in 2009 and 2012. Standard errors are shown in parentheses.

Balance at separation	Treatment cohort	Control cohort	<i>p</i> -value of difference
\$0 - \$999	0.947 (0.005)	0.913 (0.016)	0.043
\$1,000 - \$4,999	0.486 (0.016)	0.469 (0.025)	0.563
\$5,000 - \$9,999	0.488 (0.026)	0.504 (0.037)	0.738
\$10,000 - \$19,999	0.375 (0.025)	0.360 (0.037)	0.745
≥ \$20,000	0.158 (0.015)	0.111 (0.017)	0.036

Table 5. Automatic policy treatment effects on equivalent constant contribution rates

This table shows estimates of the effect of the automatic savings policy implemented at each firm on equivalent constant contribution rates using different methodologies to compute equivalent constant contribution rates. Panel A contains estimates that compare treatment to control cohorts without additional controls. Panel B contains estimates that additionally control for gender, log of salary, a quadratic in age at hire, and a spline for maximum tenure achieved (with knot points every six months) censored at 60 months. Panel C contains the control cohort's average equivalent constant contribution rate using the most comprehensive methodology. Standard errors are shown in parentheses.

Panel A: Without controls									
	Autoenrollment firms				Auto-escalation firms		Autoenrollment and escalation firms		
	Firm A	Firm B	Firm C	Firm D	Firm E	Firm F	Firm G	Firm H	Firm I
Extrapolate from 1 year, assume 85% auto-escalation take-up	--	--	--	--	2.39%	1.73%	3.56%	4.15%	2.63%
					(0.15)	(0.12)	(0.18)	(0.03)	(0.30)
Extrapolate from 1 year of tenure	1.23%	2.11%	4.41%	0.30%	1.50%	0.83%	3.19%	3.92%	2.16%
	(0.05)	(0.39)	(0.45)	(0.62)	(0.15)	(0.12)	(0.18)	(0.03)	(0.31)
Extrapolate from 5 years of tenure	0.82%	2.34%	3.64%	0.25%	1.37%	0.73%	2.36%	3.06%	1.96%
	(0.05)	(0.39)	(0.46)	(0.58)	(0.15)	(0.12)	(0.18)	(0.03)	(0.30)
Incorporate incomplete vesting of employer contributions	0.82%	2.34%	2.70%	0.09%	1.37%	0.73%	1.55%	2.46%	1.59%
	(0.05)	(0.39)	(0.45)	(0.49)	(0.15)	(0.11)	(0.16)	(0.03)	(0.30)
Incorporate withdrawals (comprehensive estimate)	0.21%	1.50%	1.46%	-0.52%	1.01%	0.45%	0.46%	1.13%	0.80%
	(0.04)	(0.32)	(0.38)	(0.38)	(0.14)	(0.10)	(0.13)	(0.03)	(0.27)
Use Vanguard's auto-escalation take-up rate	--	--	--	--	1.65%	0.94%	0.58%	1.30%	1.33%
					(0.14)	(0.10)	(0.13)	(0.03)	(0.26)

Panel B: With demographic and attrition controls									
	Autoenrollment firms				Auto-escalation firms		Autoenrollment and escalation firms		
	Firm A	Firm B	Firm C	Firm D	Firm E	Firm F	Firm G	Firm H	Firm I
Extrapolate from 1 year, assume 85% auto-escalation take-up	--	--	--	--	1.74% (0.14)	1.87% (0.11)	3.91% (0.16)	4.20% (0.04)	2.37% (0.29)
Extrapolate from 1 year of tenure	1.28% (0.04)	1.55% (0.37)	4.52% (0.48)	1.47% (0.70)	0.82% (0.14)	0.83% (0.12)	3.55% (0.16)	3.91% (0.04)	1.89% (0.30)
Extrapolate from 5 years of tenure	0.80% (0.04)	1.73% (0.37)	3.46% (0.47)	1.38% (0.66)	0.62% (0.13)	0.70% (0.11)	2.73% (0.15)	2.79% (0.04)	1.70% (0.28)
Incorporate incomplete vesting of employer contributions	0.80% (0.04)	1.73% (0.37)	2.52% (0.44)	0.87% (0.55)	0.62% (0.13)	0.58% (0.10)	1.90% (0.12)	2.32% (0.04)	1.31% (0.26)
Incorporate withdrawals (comprehensive estimate)	0.17% (0.03)	0.87% (0.28)	1.37% (0.37)	0.08% (0.43)	0.21% (0.11)	0.26% (0.08)	0.79% (0.10)	1.15% (0.03)	0.50% (0.22)
Use Vanguard's auto-escalation take-up rate	--	--	--	--	0.91% (0.11)	0.82% (0.08)	0.93% (0.09)	1.37% (0.03)	1.03% (0.21)
Panel C: Control cohort average									
Control cohort avg. equivalent constant contribution rate	1.28% (0.03)	5.27% (0.18)	3.98% (0.27)	2.82% (0.31)	7.55% (0.11)	5.37% (0.07)	2.48% (0.10)	1.69% (0.02)	7.14% (0.17)

Table 6. Average cash leakage rates by cohort × firm

This table shows the average cash leakage rate applied upon separation for each cohort within each firm. Standard errors are shown in parentheses.

	Autoenrollment firms				Auto-escalation firms		Autoenrollment and escalation firms		
	Firm A	Firm B	Firm C	Firm D	Firm E	Firm F	Firm G	Firm H	Firm I
Treatment cohort	0.668 (0.003)	0.309 (0.010)	0.380 (0.014)	0.620 (0.021)	0.328 (0.003)	0.362 (0.003)	0.636 (0.007)	0.551 (0.002)	0.303 (0.007)
Control cohort	0.494 (0.004)	0.290 (0.009)	0.265 (0.013)	0.522 (0.021)	0.332 (0.006)	0.351 (0.004)	0.468 (0.010)	0.420 (0.002)	0.245 (0.006)
Difference	0.174 (0.005)	0.019 (0.013)	0.116 (0.019)	0.098 (0.030)	-0.004 (0.006)	0.011 (0.005)	0.167 (0.012)	0.131 (0.003)	0.057 (0.009)

Appendix Table 1. Acceptance rate of default auto-escalation in broader sample of firms

This table shows, for each firm in our broader sample, the date on which the default auto-escalation policy was implemented, the initial default contribution rate, the contribution rate at which auto-escalation would cease, the fraction of active employees that increase their contribution rate by 1% of income on each of the first three escalation dates, the fraction of active employees who are missing contribution rate data in the month before or the month of each escalation date (and hence appear in neither the numerator nor the denominator of the escalation acceptance rate), and the fraction of active employees with non-missing data and contribution rates less than the escalation cap in the month before the escalation date who then increase their contribution rate by 1% of income on each of the first three escalation dates.

Firm	Policy intro date	Policy description	Escalation acceptance rate			Missing contribution rate on escalation date			Escalation acceptance rate conditional on contributing less than escalation max		
			Date 1	Date 2	Date 3	Date 1	Date 2	Date 3	Date 1	Date 2	Date 3
C	Jul 3, 2008	5% initial, escalate up to 10%	65.4%	55.6%	31.6%	0.0%	0.0%	0.0%	81.0%	65.2%	35.3%
E	Jan 1, 2011	6% initial, escalate up to 15%	30.3%	32.4%	36.8%	0.0%	1.4%	29.6%	31.3%	33.8%	37.9%
F	Jan 1, 2011	3% initial, escalate up to 10%	26.8%	29.1%	26.9%	3.0%	0.7%	0.2%	31.4%	35.1%	34.2%
G	Jan 1, 2010	2% initial, escalate up to 6%	57.1%	41.3%	28.2%	3.0%	2.7%	1.4%	70.5%	67.1%	54.9%
H	Jan 1, 2006	3% initial, escalate up to 15%	66.8%	62.3%	--	2.0%	65.2%	--	67.8%	63.5%	--
I	Jul 1, 2007	3% initial, escalate up to 6%	17.1%	10.9%	8.7%	0.1%	0.7%	1.8%	43.2%	34.0%	7.3%
J	Jul 1, 2012	6% initial, escalate up to 10%	33.4%	38.9%	34.8%	13.0%	7.7%	8.4%	38.9%	46.9%	45.4%
K	Jul 1, 2013	6% initial, escalate up to 10%	49.1%	49.2%	41.4%	0.3%	0.0%	0.0%	56.4%	59.0%	52.4%
L	Jan 1, 2014	6% initial, escalate up to 10%	37.8%	34.9%	29.4%	5.1%	1.9%	0.3%	47.4%	44.6%	43.2%
M	Mar 1, 2010	6% initial, escalate up to 10%	16.4%	9.7%	28.8%	57.6%	51.7%	22.3%	18.7%	11.3%	8.7%

N	Sept 1, 2005	3% initial, escalate up to 6%	34.7%	24.5%	18.5%	0.5%	0.5%	0.5%	63.5%	54.7%	44.6%
O	Jan 1, 2008	4% initial, escalate up to 15%	30.2%	26.3%	13.9%	0.0%	0.0%	0.0%	30.3%	26.4%	13.9%
P	Jan 1, 2006	2% initial, escalate up to 6%	47.0%	35.3%	30.7%	20.2%	6.3%	3.1%	66.1%	47.7%	54.0%
Q	Oct 1, 2013	6% initial, escalate up to 10%	43.7%	32.4%	30.4%	0.7%	0.2%	0.1%	49.4%	37.2%	37.4%
R	Jan 1, 2008	3% initial, escalate up to 6%	53.8%	45.3%	38.6%	0.0%	0.0%	0.0%	78.8%	74.0%	68.8%
S	Jul 1, 2011	3% initial, escalate up to 10%	13.6%	13.0%	12.4%	24.5%	31.9%	37.2%	15.5%	15.3%	14.8%
T	Jan 1, 2007	3% initial, escalate up to 10%	65.0%	53.5%	39.6%	12.5%	14.7%	11.4%	73.6%	60.9%	46.4%
U	Jan 1, 2008	2% initial, escalate up to 6%	36.8%	34.0%	28.3%	64.3%	63.9%	64.7%	45.9%	46.7%	41.5%
V	Apr 6, 2009	3% initial, escalate up to 5%	24.1%	20.8%	3.8%	48.0%	49.3%	51.0%	66.1%	63.9%	26.6%
W	Jan 1, 2008	2% initial, escalate up to 6%	57.3%	38.3%	28.9%	2.0%	1.8%	1.2%	71.2%	54.5%	46.0%
X	Jan 1, 2007	5% initial, escalate up to 10%	86.6%	73.1%	66.7%	52.4%	53.4%	54.0%	88.0%	76.0%	69.9%

Figure 1. Average acceptance rate of default auto-escalation

This figure shows the average auto-escalation acceptance rate on the first three escalation dates across 21 firms that implemented default auto-escalation, a broader sample of firms than is analyzed in the main text. See Appendix B for details.

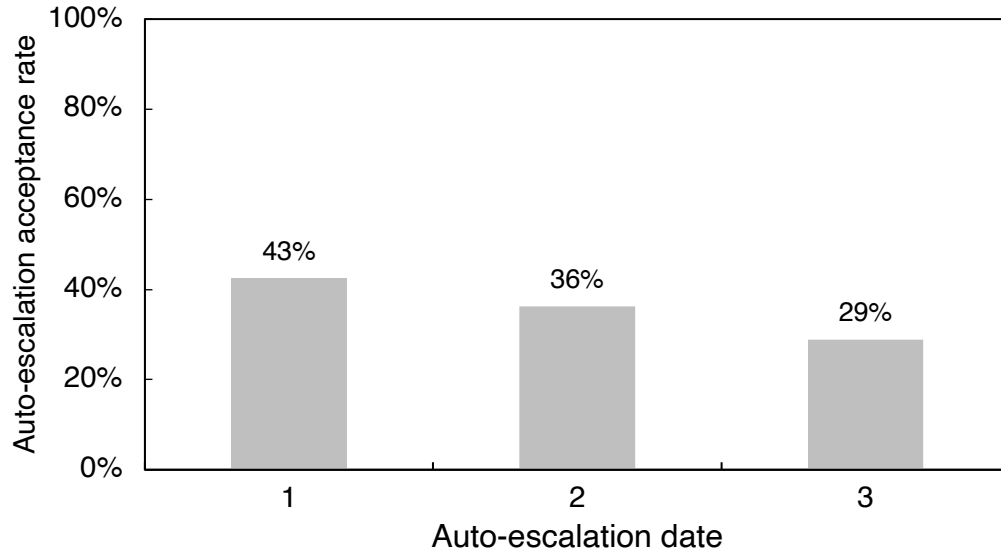


Figure 2. Cash leakage rate in year of employment separation, by 401(k) balance at separation

This figure shows the fraction of 401(k) balances at separation that are withdrawn in the calendar year of employment separation without being directly rolled over to another retirement account, by 401(k) balance at separation. The sample is treatment and control cohort employees who separated in July, August, or September at all firms except Firm H, which is missing withdrawals data in 2009 and 2012. The horizontal position of each data point indicates the center of its balance bin. Whiskers denote 95% confidence intervals.


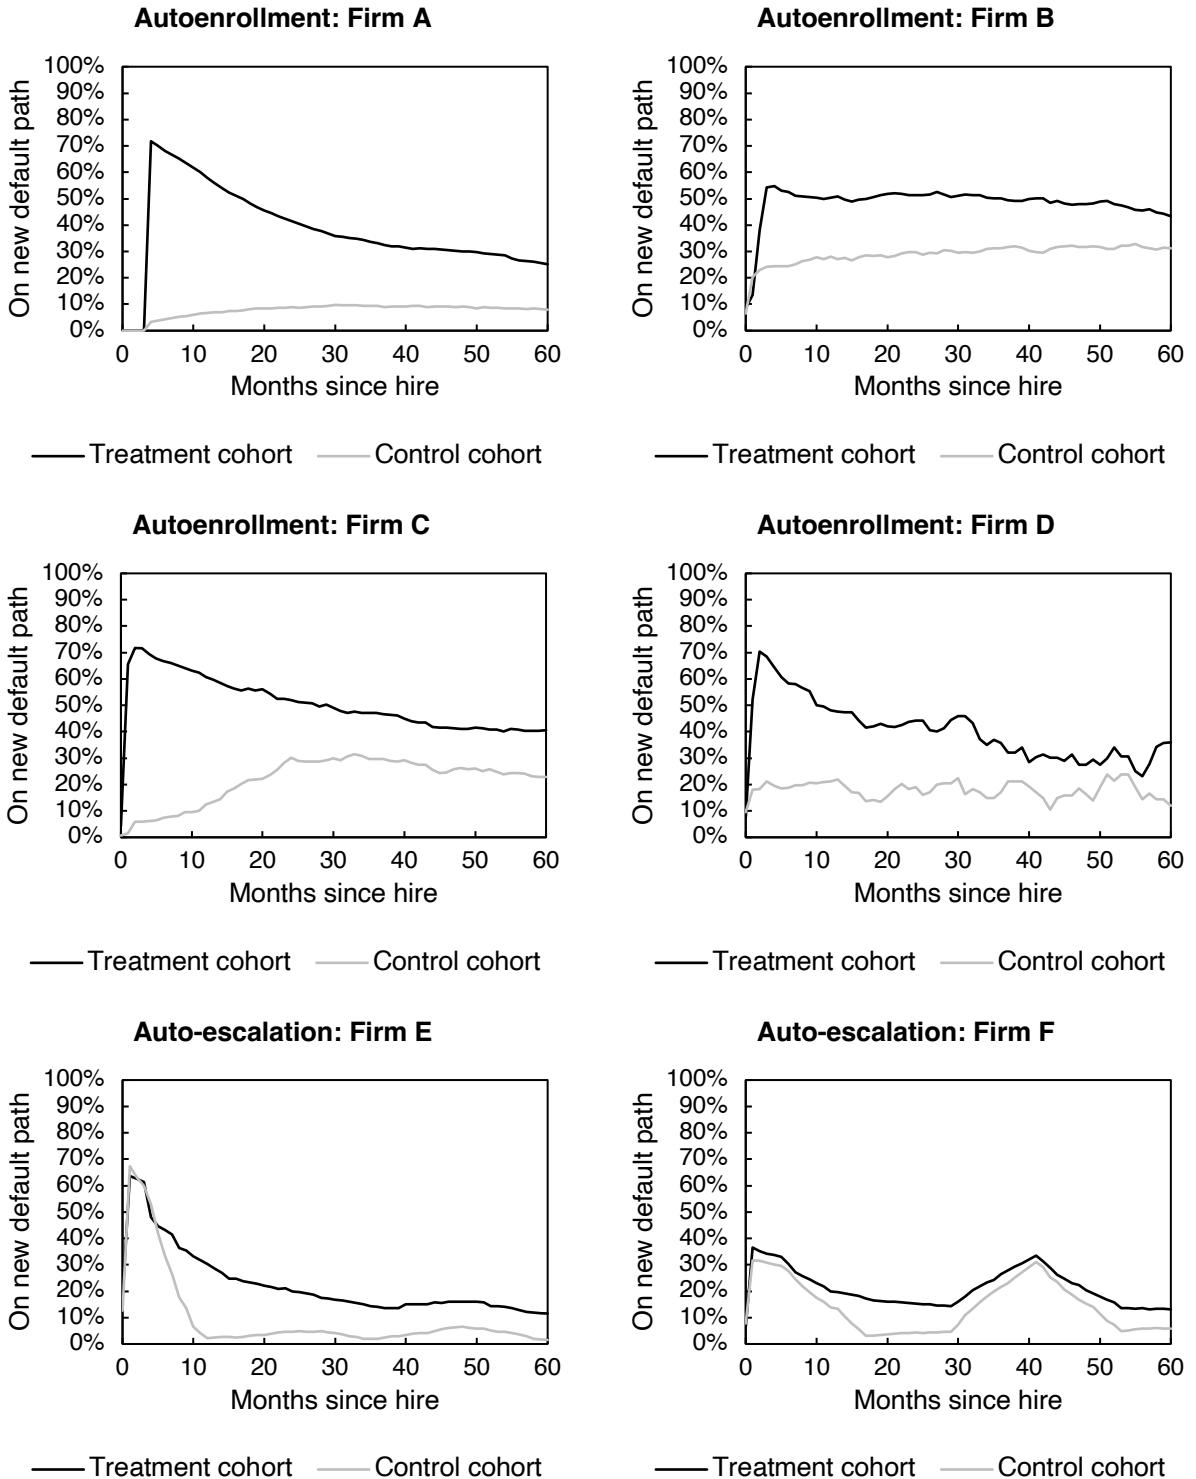


Figure 3. Percent of active employees who are on the new default contribution path

These graphs show the percent of active (i.e., still employed at the firm) employees whose contribution rate matches the contribution rate of a completely passive employee in the automatic policies regime each firm moved to. At firms with auto-escalation, the completely passive path for control employees auto-escalates at the same tenure time that auto-escalation occurs for treatment employees.



Autoenroll and escalation: Firm G



— Treatment cohort — Control cohort

Autoenroll and escalation: Firm H



— Treatment cohort — Control cohort

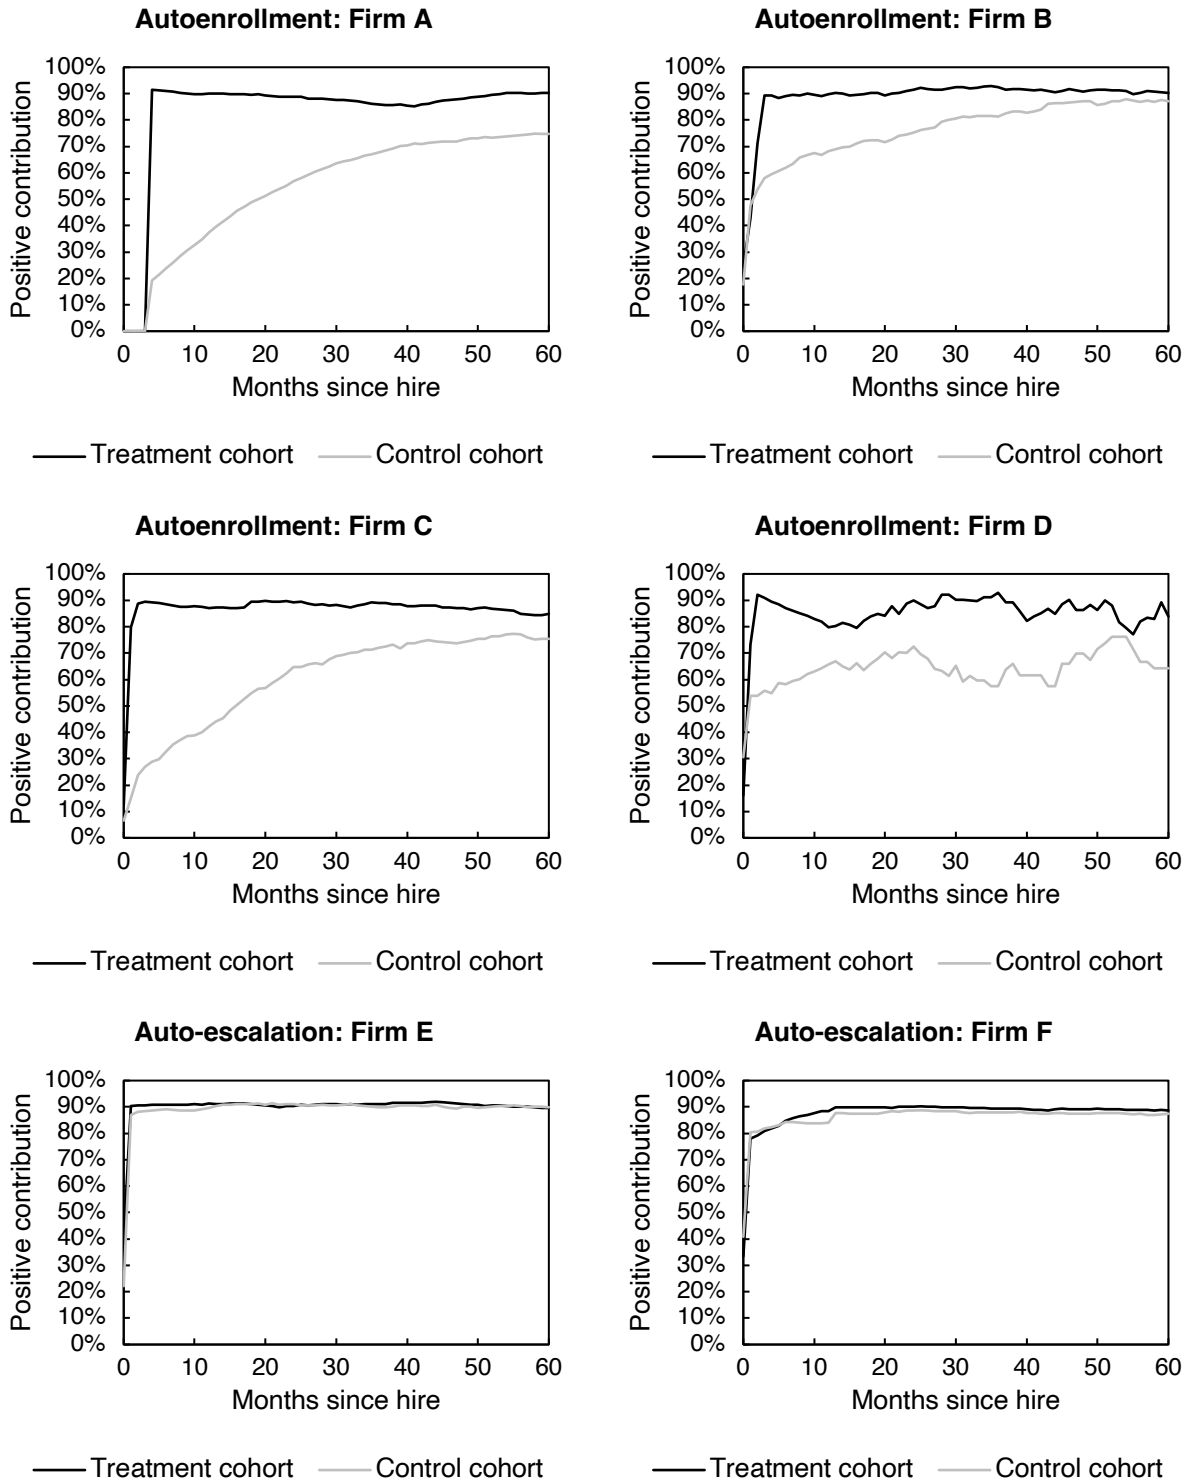
Autoenroll and escalation: Firm I



— Treatment cohort — Control cohort

Figure 4. Percent of active employees with a positive contribution rate

These graphs show the percent of active (i.e., still employed at the firm) employees whose contribution rate is positive.

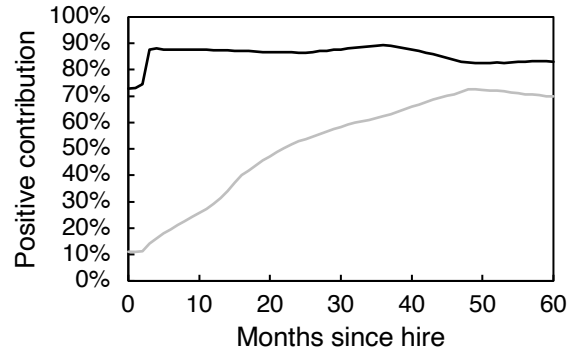


Autoenroll and escalation: Firm G



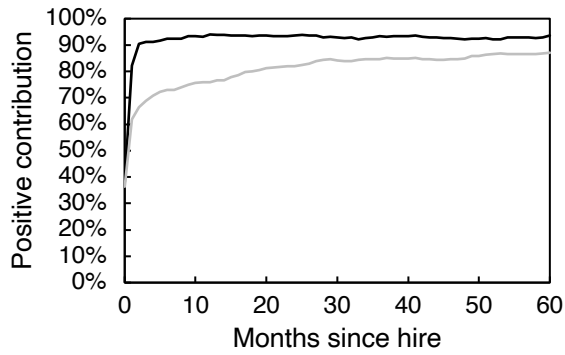
— Treatment cohort — Control cohort

Autoenroll and escalation: Firm H



— Treatment cohort — Control cohort

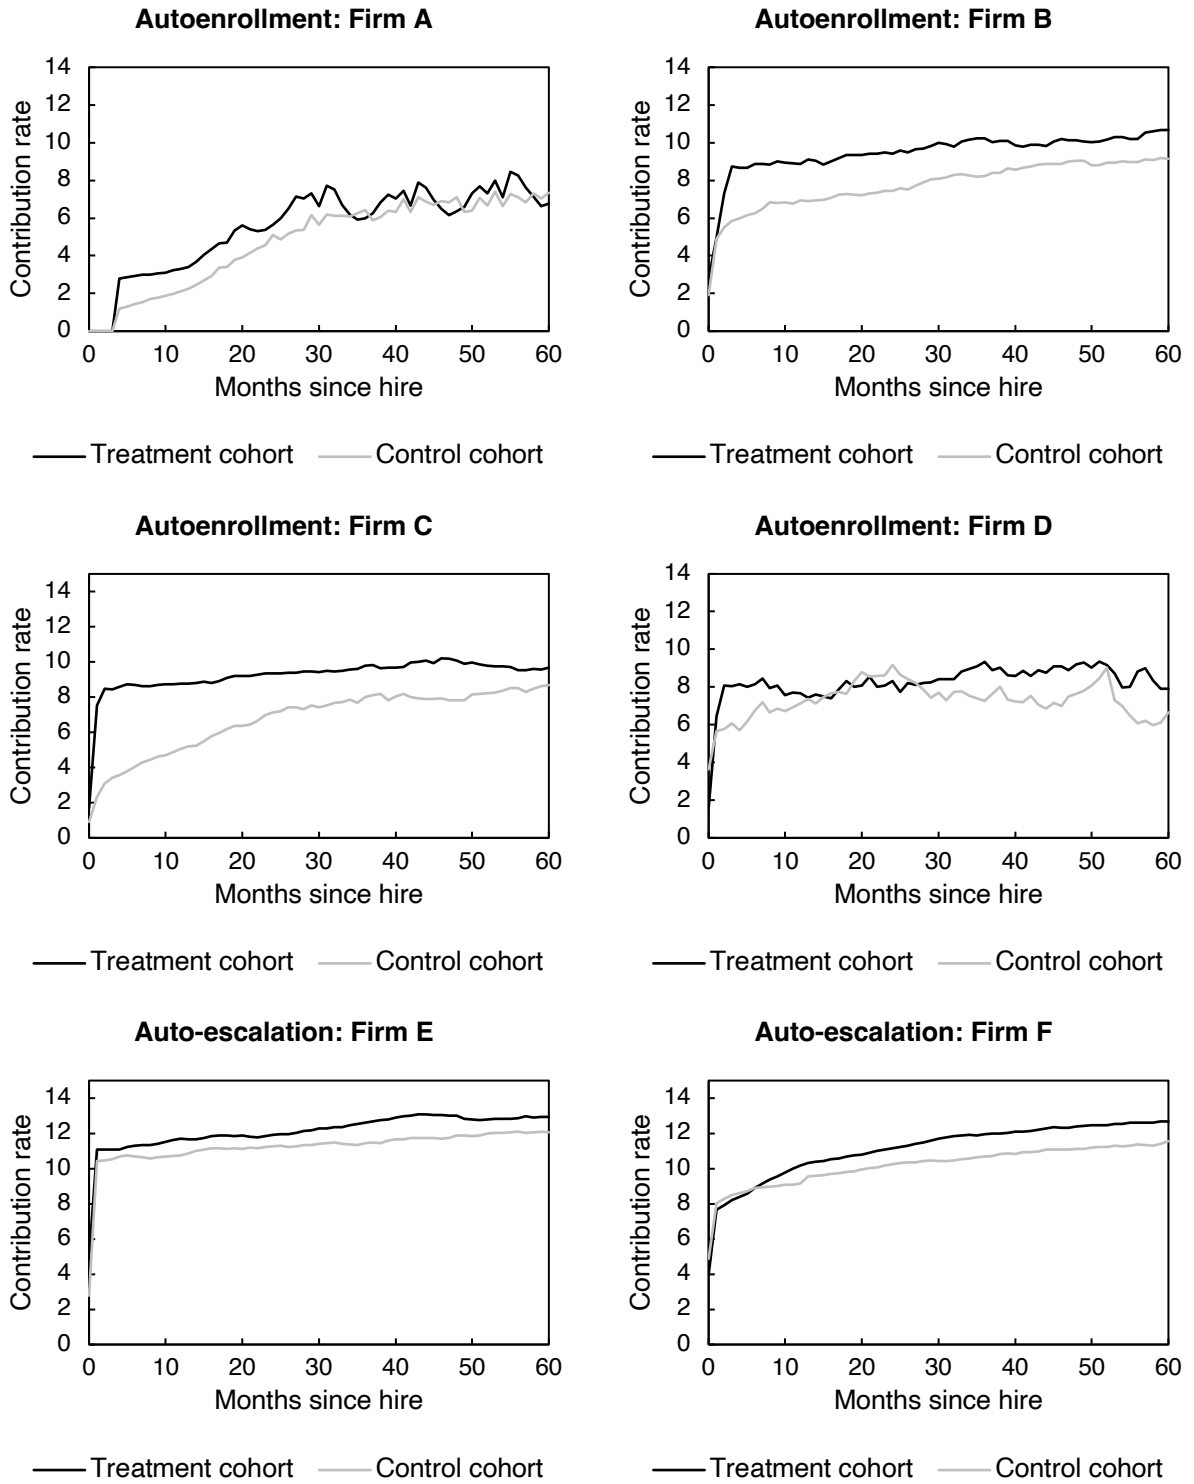
Autoenroll and escalation: Firm I



— Treatment cohort — Control cohort

Figure 5. Average employee plus match contribution rate among active employees

These graphs show the average contribution rate (employee plus match, assuming 100% vesting) of active (i.e., still employed at the firm) employees.




Autoenroll and escalation: Firm G



— Treatment cohort — Control cohort

Autoenroll and escalation: Firm H



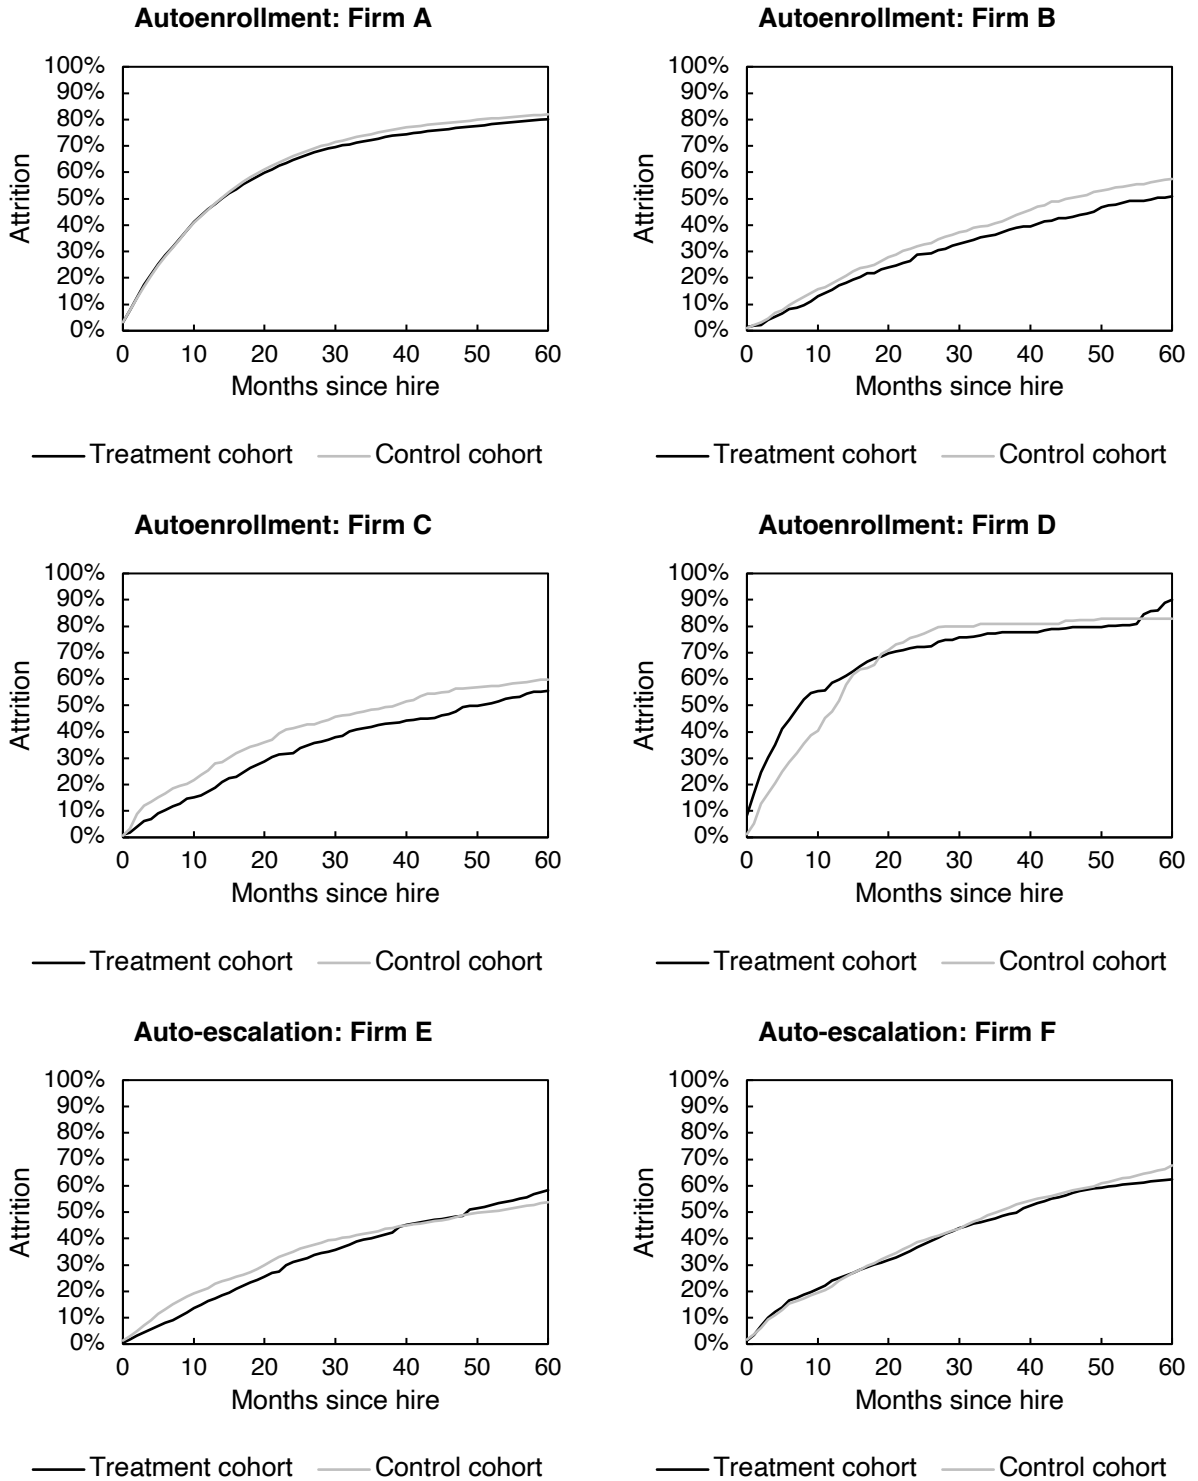
— Treatment cohort — Control cohort

Autoenroll and escalation: Firm I



— Treatment cohort — Control cohort

Figure 6. Cumulative employment attrition rates by time since hire



Autoenroll and escalate: Firm G



— Treatment cohort — Control cohort

Autoenroll and escalate: Firm H



— Treatment cohort — Control cohort

Autoenroll and escalate: Firm I



— Treatment cohort — Control cohort

Figure 7. Acceptance of auto-escalation default

These graphs show the percent of active (i.e., still employed) employees (or active employees who are contributing below the auto-escalation cap just before the auto-escalation date) at each cohort × firm who increase their contribution rate by 1% of income on their first five auto-escalation dates. At Firm H, take-up of auto-escalation is not measured on escalation dates 3 and 4 for the treatment cohort and escalation dates 4 and 5 for the control cohort because we are missing all 2009 contribution rates.

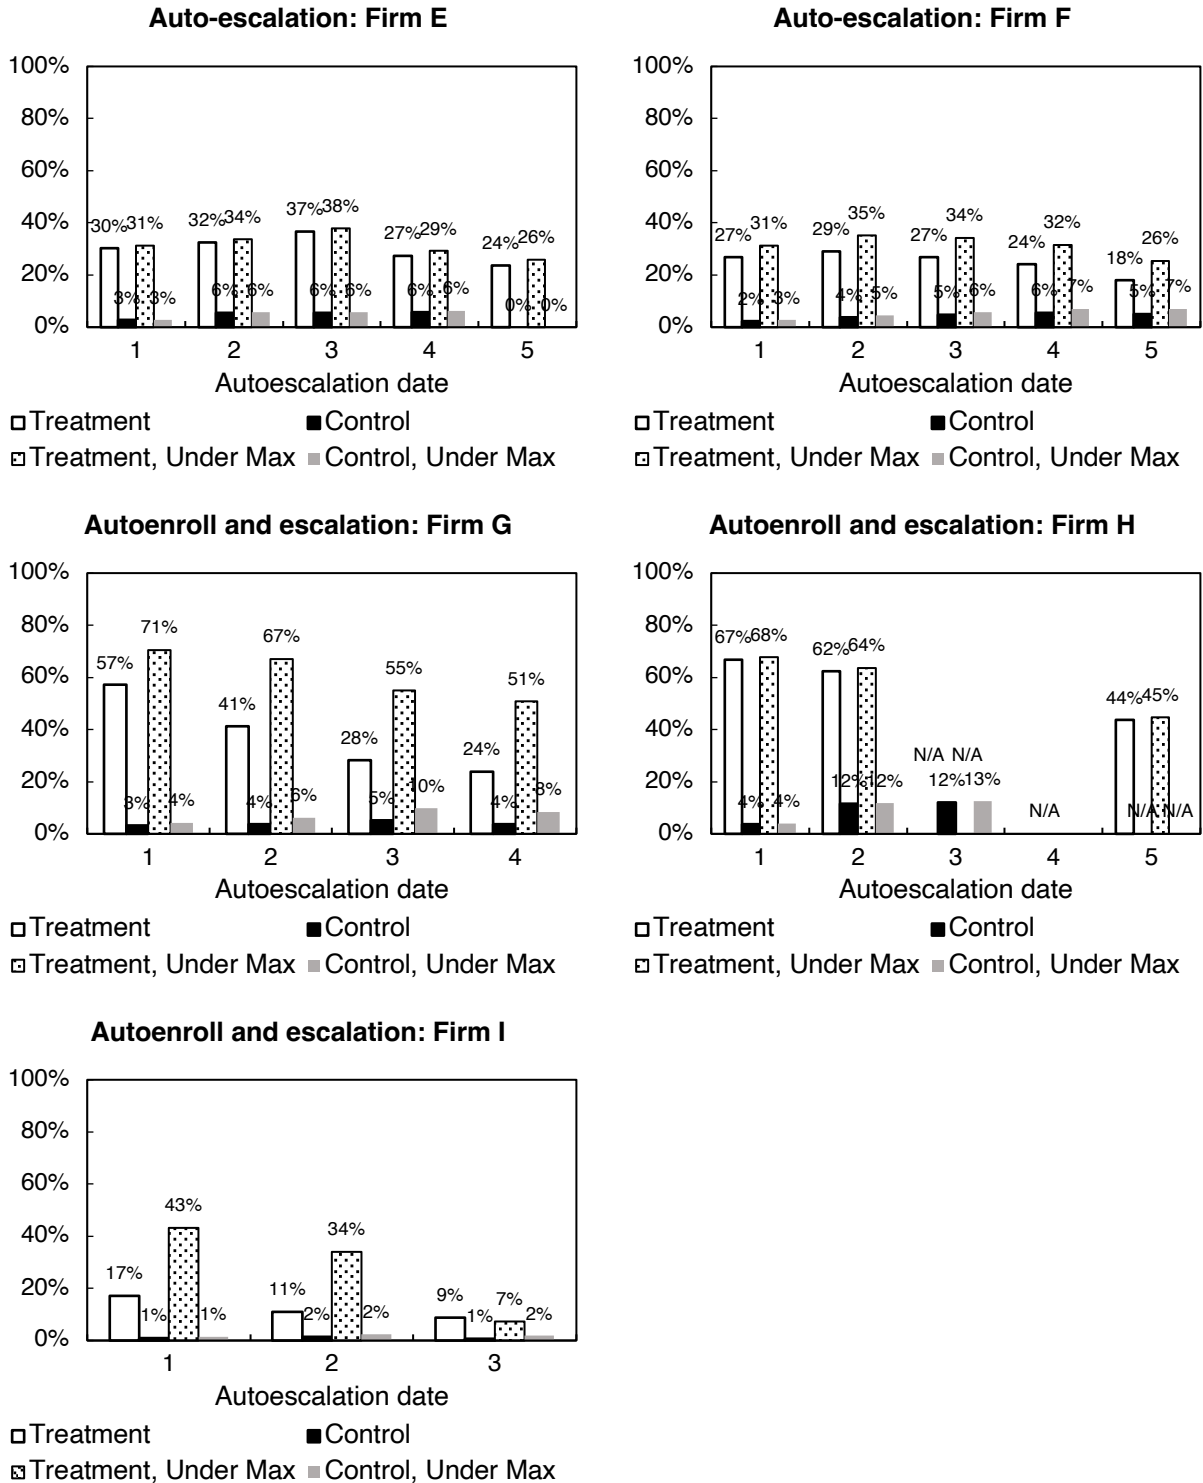
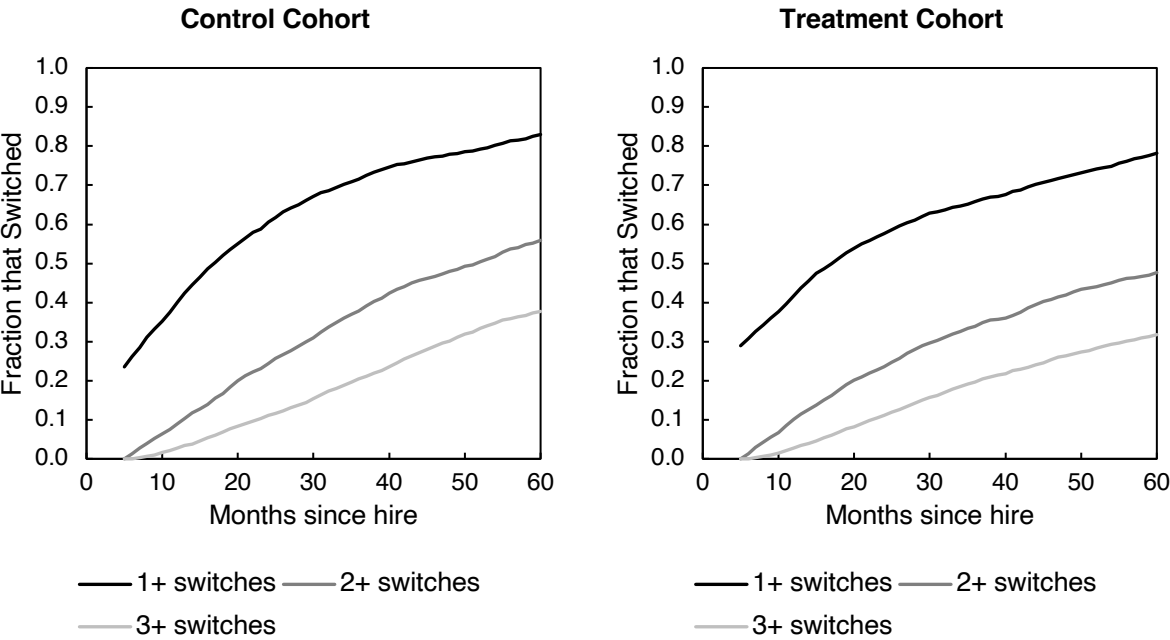


Figure 8. Contribution rate changes among active employees at automatic enrollment firms

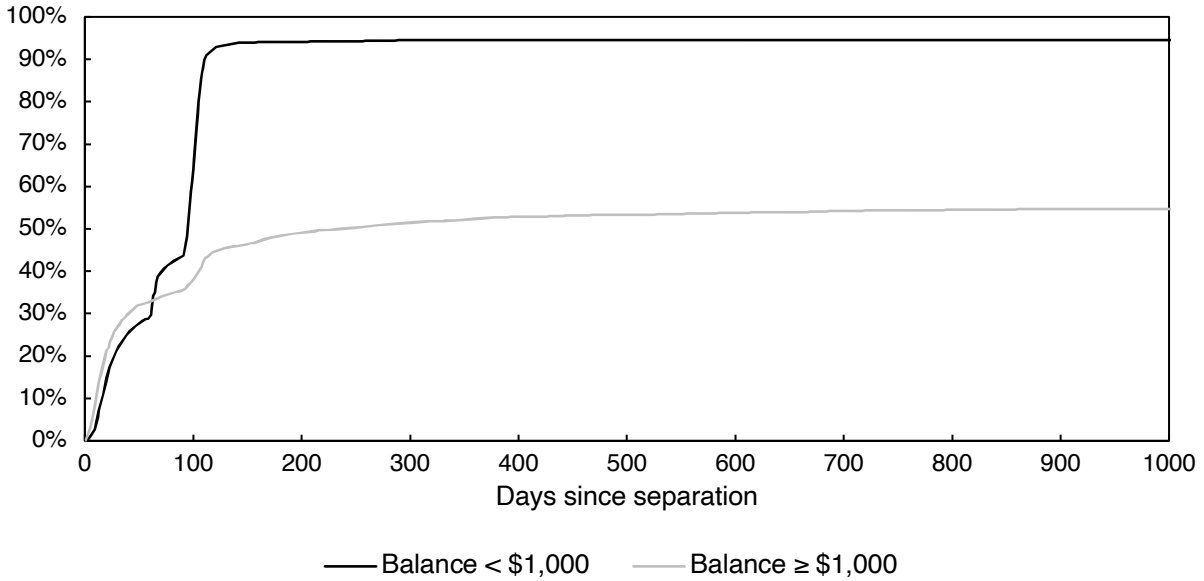
This graph shows the percent of active (i.e., still employed) employees that have actively switched their contribution rate at least once, twice, or three times by each tenure month. Rates are computed across all active employees, separately for treatment and control cohorts, at the four automatic enrollment firms. We do not begin counting contribution rate changes until tenure month 5, as we observed that some firms did not set an employee’s contribution rate to the automatic enrollment default until approximately tenure month 4.




**Appendix Figure 1. Cumulative cash leakage and rollover rate
by days since employment separation**

This figure shows, separately for those with 401(k) balances at employment separation below or above \$1,000, the cumulative cash leakage rate and cumulative rollover rate by days since separation. The sample is treatment and control cohort employees who separated in July, August, or September at all firms except Firm H, which is missing withdrawals data in 2009 and 2012.

Cumulative cash leakage rate



Cumulative rollover rate



Online Appendix Table 1. Hire cohort characteristics for balanced sample

The top two sections show hire date windows (in days relative to the firm’s automatic policy introduction), average characteristics, employment attrition rates, and employee counts in the control and treatment cohorts selected to be maximally similar to each other using the procedure in footnote 14. Salary (deflated using CPI-U to the December 31 within the control cohort’s hire window) is measured at hire when available; otherwise, it is the employee’s first observed salary if one is available, or the median first-observed salary of everybody in the firm hired within 365 days before or after the policy implementation date. The penultimate section shows *p*-values from tests of equality across the two cohorts. The final section shows the fraction of employee-months (counting only months in which the individual is employed by the sample firm) where contribution rates are imputed, the fraction of employees whose salary is imputed to be the firm-wide median salary for the purposes of computing their final synthetic 401(k) balance, and the fraction of employees whose salary we never observe.

	Autoenrollment firms				Auto-escalation firms		Autoenrollment and escalation firms		
	Firm A	Firm B	Firm C	Firm D	Firm E	Firm F	Firm G	Firm H	Firm I
<i>Control cohort</i>									
Hire dates	[-365, -276]	[-90, -1]	[-365, -276]	[-180, -1]	[-180, -1]	[-365, -276]	[-365, -276]	[-365, -276]	[-365, -276]
Age at hire	29.7	31.5	36.0	35.4	35.6	37.5	32.8	35.5	32.3
Female	65.5%	45.3%	22.9%	32.5%	20.4%	22.5%	64.0%	77.8%	41.6%
Salary	\$23,104	\$55,006	\$62,019	\$46,641	\$40,290	\$60,915	\$33,580	\$32,149	\$60,205
5-year attrition	83.8%	60.9%	50.0%	86.7%	51.9%	72.6%	77.9%	70.6%	50.0%
Employees	3,529	161	96	120	1,293	1,167	467	7,483	466
<i>Treatment cohort</i>									
Hire dates	[0, 89]	[0, 89]	[0, 89]	[0, 179]	[0, 179]	[0, 89]	[0, 89]	[0, 89]	[0, 89]
Age at hire	29.6	30.9	37.2	35.0	34.8	37.6	33.2	35.5	32.4
Female	66.3%	41.8%	23.3%	36.3%	19.1%	21.4%	61.6%	78.3%	40.9%
Salary	\$22,975	\$56,295	\$61,888	\$42,528	\$39,732	\$63,028	\$35,128	\$34,998	\$58,405
5-year attrition	81.1%	54.1%	60.0%	82.3%	53.8%	63.8%	78.3%	71.4%	50.4%
Employees	3,763	146	90	124	1,761	1,750	479	8,320	403
<i>Tests of equality of means and proportions between cohorts</i>									
Age at hire	0.926	0.551	0.437	0.740	0.062	0.908	0.540	0.974	0.979
Female	0.493	0.531	0.947	0.535	0.360	0.480	0.438	0.504	0.837
Salary	0.787	0.784	0.984	0.362	0.561	0.214	0.541	< 0.001	0.456
5-year attrition	0.002	0.233	0.172	0.344	0.289	< 0.001	0.898	0.305	0.913
<i>Imputed or missing observations</i>									
Contrib. imputed	9.7%	0.4%	5.0%	0.0%	1.5%	2.9%	1.5%	25.2%	0.5%
Salary imputed	0.1%	0.0%	3.2%	17.6%	7.2%	12.2%	0.0%	16.3%	1.2%
Salary never observed	3.0%	0.0%	14.5%	33.6%	9.7%	21.2%	1.6%	37.3%	4.1%

Online Appendix Table 2. Automatic policy treatment effects on equivalent constant contribution rates for balanced sample

This table shows estimates of the effect of the automatic savings policy implemented at each firm on equivalent constant contribution rates using different methodologies to compute equivalent constant contribution rates. The control and treatment cohorts are selected to be maximally similar to each other using the procedure in footnote 14. Panel A contains estimates that compare treatment to control cohorts without additional controls. Panel B contains estimates that additionally control for gender, log of salary, a quadratic in age at hire, and a spline for maximum tenure achieved (with knot points every six months) censored at 60 months. The final row contains the control cohort’s average equivalent constant contribution rate using the most comprehensive methodology. Standard errors are shown in parentheses.

Panel A: Without controls									
	Autoenrollment firms				Auto-escalation firms		Autoenrollment and escalation firms		
	Firm A	Firm B	Firm C	Firm D	Firm E	Firm F	Firm G	Firm H	Firm I
Extrapolate from 1 year, assume 85% auto-escalation take-up	--	--	--	--	1.54% (0.19)	1.87% (0.23)	3.69% (0.37)	4.34% (0.07)	2.22% (0.47)
Extrapolate from 1 year of tenure	1.32% (0.09)	1.62% (0.79)	3.89% (1.03)	1.48% (0.80)	0.92% (0.20)	1.17% (0.23)	3.46% (0.37)	4.15% (0.07)	1.76% (0.48)
Extrapolate from 5 years of tenure	0.94% (0.09)	1.74% (0.73)	2.95% (1.04)	1.32% (0.78)	0.58% (0.19)	0.88% (0.23)	2.83% (0.38)	3.10% (0.07)	1.98% (0.47)
Incorporate incomplete vesting of employer contributions	0.94% (0.09)	1.74% (0.73)	1.65% (1.02)	0.88% (0.63)	0.58% (0.19)	0.86% (0.22)	2.13% (0.35)	2.47% (0.07)	1.48% (0.45)
Incorporate withdrawals (comprehensive estimate)	0.26% (0.07)	1.09% (0.58)	0.48% (0.90)	0.32% (0.42)	0.13% (0.18)	0.59% (0.19)	0.81% (0.29)	1.10% (0.06)	0.70% (0.40)
Use Vanguard’s auto-escalation take-up rate	--	--	--	--	0.55% (0.18)	1.08% (0.20)	0.94% (0.29)	1.18% (0.06)	1.12% (0.40)

Panel B: With demographic and attrition controls									
	Autoenrollment firms				Auto-escalation firms		Autoenrollment and escalation firms		
	Firm A	Firm B	Firm C	Firm D	Firm E	Firm F	Firm G	Firm H	Firm I
Extrapolate from 1 year, assume 85% auto-escalation take-up	--	--	--	--	1.78% (0.18)	2.24% (0.23)	3.95% (0.34)	4.78% (0.08)	2.11% (0.43)
Extrapolate from 1 year of tenure	1.33% (0.08)	1.39% (0.79)	3.22% (1.20)	2.57% (0.83)	1.06% (0.20)	1.38% (0.24)	3.71% (0.35)	4.50% (0.09)	1.63% (0.45)
Extrapolate from 5 years of tenure	0.87% (0.07)	1.41% (0.72)	1.88% (1.15)	2.36% (0.77)	0.67% (0.18)	0.99% (0.24)	3.02% (0.34)	3.00% (0.08)	1.87% (0.43)
Incorporate incomplete vesting of employer contributions	0.87% (0.07)	1.41% (0.72)	1.06% (1.09)	1.70% (0.62)	0.67% (0.18)	0.94% (0.21)	2.25% (0.29)	2.57% (0.08)	1.42% (0.39)
Comprehensive estimate	0.17% (0.06)	0.61% (0.52)	0.17% (0.95)	0.86% (0.45)	0.27% (0.15)	0.62% (0.18)	1.01% (0.22)	1.42% (0.06)	0.63% (0.33)
Use Vanguard's auto-escalation take-up rate	--	--	--	--	0.77% (0.14)	1.23% (0.18)	1.15% (0.21)	1.54% (0.06)	1.07% (0.32)
Panel C: Control cohort average									
Control cohort avg. equivalent constant contribution rate	1.19% (0.05)	5.22% (0.40)	5.27% (0.71)	2.10% (0.29)	7.70% (0.14)	4.95% (0.15)	2.50% (0.19)	1.72% (0.04)	6.56% (0.30)

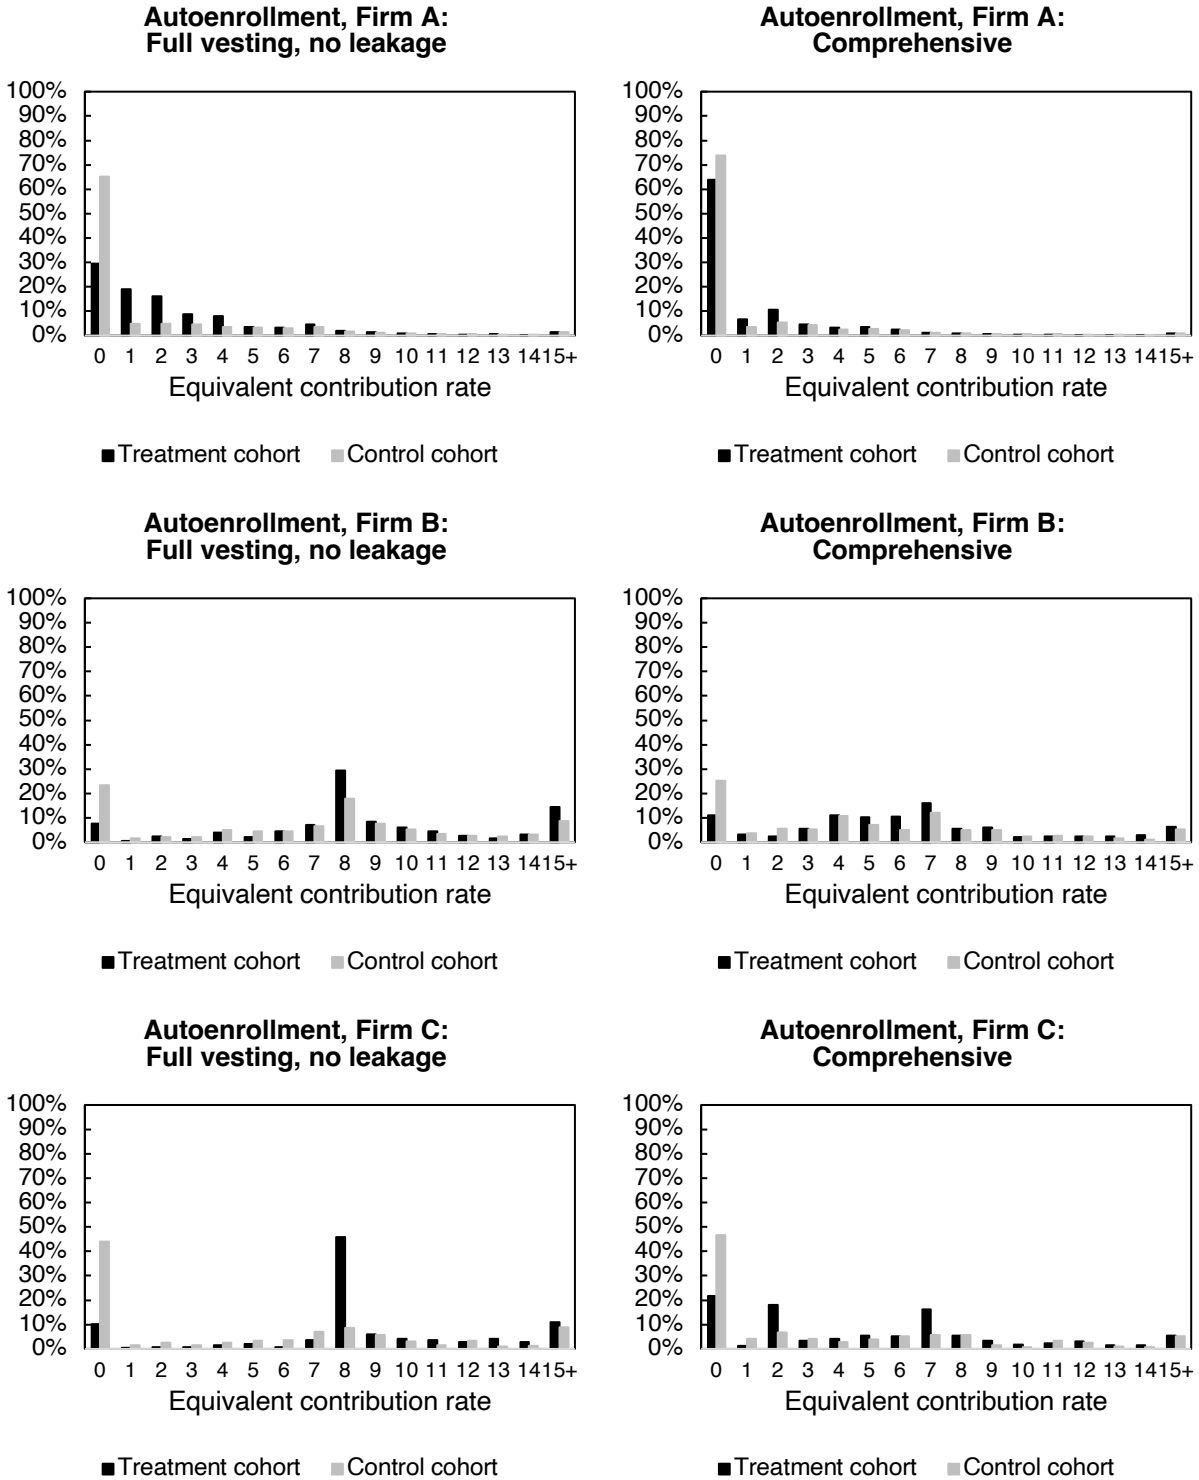
**Online Appendix Table 3. Cash leakage rates by age
and 401(k) balance at employment separation**

This figure shows the cash leakage rate in the year of employment separation by 401(k) balance and age at separation. The sample is treatment and control cohort employees who separated in July, August, or September at all firms except Firm H, which is missing withdrawals data in 2009 and 2012. Standard errors are shown in parentheses.


<u>Balance at separation</u>	<u>Age</u>				
	<u>0 – 24</u>	<u>25 – 29</u>	<u>30 – 34</u>	<u>35 – 44</u>	<u>≥ 45</u>
\$0 - \$999	0.942 (0.006)	0.950 (0.010)	0.947 (0.014)	0.923 (0.015)	0.957 (0.014)
\$1,000 - \$4,999	0.457 (0.022)	0.554 (0.030)	0.512 (0.040)	0.444 (0.031)	0.465 (0.034)
\$5,000 - \$9,999	0.355 (0.041)	0.444 (0.048)	0.574 (0.056)	0.607 (0.042)	0.510 (0.049)
\$10,000 - \$19,999	0.298 (0.041)	0.415 (0.046)	0.455 (0.055)	0.385 (0.051)	0.334 (0.042)
≥ \$20,000	0.103 (0.024)	0.128 (0.027)	0.172 (0.034)	0.172 (0.024)	0.128 (0.020)

Online Appendix Figure 1. Distributions of equivalent constant contribution rates


These graphs show the density of equivalent constant contribution rates of all employees in the analysis sample at each firm. Fractional equivalent constant contribution rates are rounded to nearest integer for this figure only. Equivalent constant contribution rates of 14.5% and above are pooled in the “15+” bin.




**Autoenrollment, Firm D:
Full vesting, no leakage**




**Autoenrollment, Firm D:
Comprehensive**




**Auto-escalation, Firm E:
Full vesting, no leakage**




**Auto-escalation, Firm E:
Comprehensive**




**Auto-escalation, Firm F:
Full vesting, no leakage**




**Auto-escalation, Firm F:
Comprehensive**




**Autoenroll and escalation, Firm G:
Full vesting, no leakage**




**Autoenroll and escalation, Firm G:
Comprehensive**



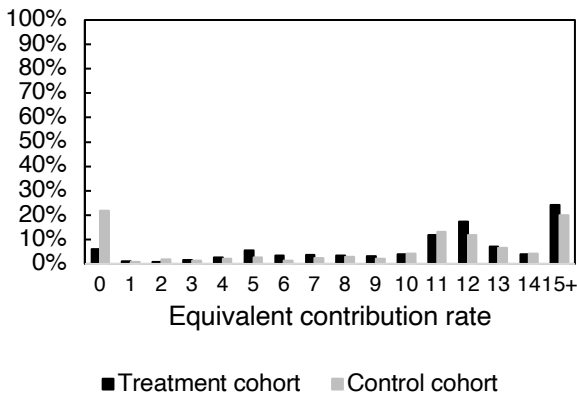
**Autoenroll and escalation, Firm H:
Full vesting, no leakage**



**Autoenroll and escalation, Firm H:
Comprehensive**



**Autoenroll and escalation, Firm I:
Full vesting, no leakage**



**Autoenroll and escalation, Firm I:
Comprehensive**

