Retention or regressivity? The empirical effects of 401(k) vesting schedules*

Guillermo Carranza
Yale University

Aaron Goodman
Vanguard

October 2025

Abstract

Vesting requirements are a common yet understudied feature of defined-contribution retirement plans. Using administrative recordkeeping data, we find that 30% of separations occur during participants' vesting periods. The resulting forfeitures of employer contributions are concentrated among lower-income participants and make the distribution of 401(k) compensation significantly more regressive. Firms do not enjoy offsetting efficiency benefits: employing both cross-plan and within-plan identification strategies, we find no evidence that vesting exerts a causal retention effect. A linked survey shows informational frictions to be a key mechanism, as a majority of respondents do not know their current plan's vesting rules.

*We thank Joseph Altonji, Bronson Argyle, Ben Barasky, John Beshears, Marcos Catao, James Choi, Taha Choukhmane, Allison Cole, Fiona Greig, Kelly Hahn, David Laibson, Cormac O'Dea, Lawrence Schmidt, Seth Zimmerman, colleagues at Vanguard's Investment Strategy Group and Strategic Retirement Consulting team, and numerous seminar participants for valuable feedback. The Vanguard Group, Inc. implemented the survey analyzed in this paper; the survey design benefited from the collaboration and assistance of Paulo Costa, Anna Madamba, and Fu Tan. The views expressed here are those of the authors and not necessarily those of the Vanguard Group, Inc. **This material is provided for informational purposes only and is not intended to be investment advice or a recommendation to take any particular investment action.**

1 Introduction

Employer contributions to 401(k) plans are an essential feature of the defined contribution retirement system. Now totaling over $200 billion per year,[1] 401(k) compensation represents a substantial cost for firms and an important source of retirement wealth for workers. An underappreciated fact about these large contribution flows is that some are eventually reversed. At plans with vesting requirements, workers do not have immediate ownership rights to employer contributions and must forfeit a portion of their balance if they separate before the end of their vesting period. In 2022, 50% of Vanguard-administered plans attached a vesting schedule to their matching contributions.[2]

Despite the prevalence of 401(k) vesting requirements, research assessing their empirical effects is limited. We use administrative recordkeeping data from Vanguard to study two central questions. First, what are the mechanical effects of forfeitures on workers' retirement wealth accumulation and the distribution of 401(k) compensation? Second, does vesting provide retention benefits for firms by inducing behavioral responses that lengthen employment spells?

In studying forfeitures and their wealth effects, we consider the universe of participant separations across 1,500 plans during the 2010-2022 period. For each of the 4.7 million separations in our sample, we observe the participant's job tenure, the participant's final account balance, and the amount of any employer contributions the participant forfeits due to incomplete vesting. We find that forfeitures are common events, occurring in roughly 30% of separations in the most recent annual cohort. The forfeiture rate has risen over time, in part because typical job spells have become shorter: within our balanced panel of plans, median tenure at separation falls from four years to two years between 2010 and 2022.[3]

We find that forfeiture incidence is highly heterogeneous, with losses concentrated

[1] See Table A4 in the Department of Labor's 2021 summary of Form 5500 filings.

[2] See Clark (2023). In our empirical sample, about 75% of plans make at least one type of employer contribution (matching or non-matching) that is subject to vesting requirements.

[3] Studies using data from the Current Population Survey (e.g, Copeland 2023) generally show median employment tenure declining more slowly than in our sample. Aside from possible compositional differences between our sample and the broader U.S. workforce, methodological differences also matter: we measure tenure *at separation* rather than among current employees.

among younger and lower-income participants. Because they occur most often among lower-income workers with shorter tenures, forfeitures yield modest amounts for employers but represent substantial amounts for the affected participants: conditional on being nonzero, the average forfeiture is about 40% of the participant's final account balance. Forfeiture losses become even more consequential when accounting for foregone market returns and the fact that participants may suffer multiple forfeitures during the course of their careers. In a calibrated lifetime simulation model, we find that forfeitures reduce a regular 401(k) contributor's retirement wealth by an average of $26,000 at age 65, with lifetime costs borne disproportionately by lower-income workers.

To quantify the distributional implications of vesting for retirement compensation, we compare gross employer contributions (gross flows into participants' 401(k) accounts during their employment spells) to net employer contributions (gross flows minus any forfeitures that occur at separation) across the income spectrum. Gross contributions are themselves regressive, as higher-income participants are more likely to maximize employer matches (Greig et al., 2024). But the disparity in forfeiture rates widens the net contribution gap between the top and bottom and quintiles by 0.6 percentage points of income, roughly doubling the 0.65-percentage-point gap that exists in gross terms.

Notwithstanding their nominal costs on wealth accumulation, vesting requirements could still produce efficiency benefits for firms and workers. If vesting schedules elicit behavioral responses that lengthen employment spells, the resulting cost savings (from fewer job-posting and training expenses) and productivity gains (from firm-specific human capital accumulation) would be important to weigh against mechanical forfeiture losses. Our second set of empirical analyses test whether vesting schedules exert a causal retention effect. Using a subset of participants who begin employment spells during the 2010-2022 period, we study the retention question with two different identification strategies and in each case find a null result: vesting schedules do not appear to have any significant effect on participants' separation decisions or ultimate employment tenures.

Our first test for retention effects relies on the discontinuous nature of 401(k) vesting schedules, which are stated in annual terms. Rather than accruing continuously with

each successive day or hour of employment, participants' ownership share in employer contributions jumps discretely at annual vesting deadlines (after one year of employment, then again after two years of employment, etc.). The discrete jumps in ownership share mean that the financial return to participants of additional employment service is much higher just before than just after annual vesting deadlines. If participants respond to vesting incentives, then separation probabilities should show corresponding jumps just after vesting deadlines – participants who are employed near an annual deadline should "wait to vest" before quitting. An empirical confound is the fact that other compensation-related events tend to happen at annual employment anniversaries, including revisions to base salary, payments of bonus salary, promotion decisions, and grants of non-401(k) deferred compensation that may be subject to unobservable vesting schedules. Simple estimates that compared separation probabilities before and after vesting deadlines could be biased by these confounding factors.

To isolate the retention effect of 401(k) vesting schedules, we use plans without vesting requirements as a control group. Applying a standard difference-in-difference setup, we test whether the jump in separation rates after annual employment anniversaries is larger at vesting plans than at non-vesting plans. The identification assumption – bolstered by the parallel pre-trends we observe – is that non-vesting plans capture the counterfactual separation dynamics around employment anniversaries that would have occurred absent the vesting requirements. We find no differential jump in separation probabilities and obtain difference-in-difference estimates that are consistently zero. The null result holds across the participant age and income distributions, for participants with small and large balances, and across plans with vesting schedules of different lengths.

Our second retention test relies on a different kind of discontinuity that applies at a small subset of plans in our sample. While most plans use elapsed time since hire to credit vested employment service (we restrict to these plans when applying the difference-in-difference setup), we identify 29 plans that credit service with the "equivalency method." At the 29 plans we study, the practical effect of equivalence rules is to award a year of vesting credit to participants who are continuously employed for five months and at least

one day of a sixth month in a given *calendar* year. As a result, participants hired in July or earlier receive credit for their first calendar year of employment but those hired in August or later do not. The July-August comparison thus gives a cleanly identified, within-plan estimate of the retention effect: August hires begin employment at essentially the same time, receive the same non-401(k) compensation in expectation, and are subject to the same non-vesting 401(k) rules, but face effective vesting schedules that are one year longer. If vesting exerts a retention effect, then August hires' longer vesting schedules should cause them to have longer average employment tenures.

The July-August discontinuity yields another null result. We estimate models of survival probability that control for plan fixed effects and find that August hires are not more likely to remain employed beyond given time thresholds than July hires. We view the survival estimates at equivalency-method plans as a distinct and complementary piece of evidence supporting our conclusion that vesting schedules do not exert retention effects. While the difference-in-difference approach has a larger sample and uses a more salient source of variation, the equivalency-method sample makes within-plan comparisons that control for unobserved features of firms' compensation policies. Both research designs corroborate our null retention result.

We consider two potential explanations for our null retention finding. First, workers may not respond to vesting incentives if the wage gains they realize from switching jobs outweigh their forfeiture losses. But using a subset of job switchers whose old and new firms both have Vanguard-administered 401(k) plans, we find that workers with forfeitures do not earn larger wage gains than observably similar workers who are fully vested at separation. The lack of empirical relationship between forfeitures and wage gains motivates our second hypothesis: informational frictions. Using data from a 2024 survey of Vanguard 401(k) participants, we analyze a series of questions assessing participants' vesting knowledge and check respondents' answers against the vesting rules at their plan. Information gaps appear to be an important explanation for the lack of behavioral response to vesting requirements, as most survey respondents are uninformed about vesting. In a sample of 1,018 respondents, only 31% can identify the definition of a vesting schedule from

a multiple-choice list. After being provided with the definition, only 33% of respondents can correctly state whether their plan has a vesting schedule. Information gaps are largest among the lower-income and shorter-tenured participants for whom vesting requirements are most binding.

Our findings have important policy implications. Even though employers often cite retention as a motivation for using vesting schedules – surveys of plan sponsors indicate that a substantial share view vesting as a way to reduce turnover[4] – our null results suggest that retention effects are not strong enough to rationalize current vesting rules. Rather than serving as a retention tool, vesting requirements appear to function mainly as a cost-control mechanism: many employers value their ability to use the forfeitures recouped from shorter-tenured employees to help cover the cost of 401(k) contributions for longer-tenured employees.[5] Given the correlation between tenure and income, these cost savings are realized mostly at the expense of lower-income participants and make the distribution of retirement compensation more regressive.

Policymakers have historically sought to prevent income-based discrimination in 401(k) plans and ensure that retirement benefits are equitably distributed among employees. Indeed, under the Employee Retirement Income Security Act, plan sponsors must annually demonstrate their compliance with explicit nondiscrimination criteria and show that their 401(k) benefits are not too skewed toward high-income employees. But because they only consider gross flows of employer contributions, ERISA nondiscrimination tests neglect the regressive effects of vesting requirements.[6] One could imagine nondiscrimination criteria stated in terms of net rather than gross employer contributions, though statements about optimal regulation would need to account for potential general-equilibrium effects on labor demand, wages, and other non-wage benefits.

Vesting schedules are an example of deferred-compensation contracts that firms may use to finance and retain skilled human capital (Sun and Xiaolan, 2019; Hoffmann and

[4]See, e.g., Jeszeck et al. (2016).

[5]See again Jeszeck et al. (2016). See also Ward (2022).

[6]Nondiscrimination tests only apply to plans that do not use safe-harbor designs. The safe-harbor standards themselves constrain vesting: plans without automatic enrollment must vest safe-harbor contributions immediately and those with automatic enrollment must vest safe-harbor contributions within two years.

Vladimirov, 2025). A well-developed literature shows that unvested stock options and other deferred equity compensation generally deliver strong employee retention effects (Oyer and Schaefer, 2005; Balsam et al., 2007; Aldatmaz et al., 2018; Jochem et al., 2018; Ladika and Sautner, 2020; Gong et al., 2023). Although our difference-in-difference setup is similar to the empirical approaches employed in the equity-compensation literature, we find a divergent null retention result for 401(k) vesting. The disparity likely arises from differences in monetary stakes and employee knowledge: equity compensation is usually larger in dollar terms than 401(k) contributions and managers are likely better informed about vesting rules than the plan participants in our survey. Of course, the population of workers receiving 401(k) contributions is much larger than the one receiving stock options, and our results indicate that the most prevalent form of compensation vesting among the U.S. workforce does not serve a systematic retention purpose. The primary effect of 401(k) vesting is instead direct cost containment among higher-turnover workers whose separation decisions appear inelastic to 401(k) rules.

Our work also relates to the extensive literature on the architecture of public and private retirement systems. The part of this literature studying vesting requirements in modern defined-contribution plans is limited. Choi et al. (2024) show that vesting forfeitures attenuate the net saving effects of automatic-enrollment policies, though they work with a smaller recordkeeping sample and do not estimate retention effects. Prince et al. (2024) use recent public regulatory filings to document the number of workers affected by 401(k) vesting requirements but do not explore participant-level wealth or retention effects. Most prior work examining vesting requirements has done so within the context of defined-benefit pension systems. Kotlikoff and Wise (1987) and Schiller and Weiss (1979) discuss the labor-supply incentives created by discontinuous pension vesting rules and provide early empirical evidence that workers respond to these incentives when making retirement decisions. Ni and Podgursky (2016) and Biasi (2019, 2024) replicate this finding with more recent data, showing that public school teachers' retirement decisions respond to changes in state pension systems.[7] By using data from defined-contribution plans, we study vesting

[7]See also Kim and Wroblewski (2025), who show that the federal government's shift from defined-benefit to defined-contribution retirement compensation reduced employee retention.

requirements that currently apply to a much broader share of the labor force and, though they often arose as remnants of the defined-benefit system, do not exert the retention effects these prior studies document.

Finally, we add to a relatively recent strand of the literature focusing on inequality in the 401(k) system. National studies have shown substantial gaps in retirement readiness between high- and low-income workers (Tan et al., 2023; Sabelhaus and Volz, 2022). Similar retirement wealth gaps exist between racial groups (Hou and Sanzenbacher, 2020; Wolff, 2023; Choukhmane et al., 2023). In explaining these gaps, the literature has documented income and racial disparities in 401(k) participation rates, in saving rates conditional on participation, and in leakage arising from pre-retirement withdrawals (Choukhmane et al., 2023). We identify vesting requirements as another institutional factor that perpetuates retirement wealth inequality: even among workers who participate in 401(k) plans and have similar contribution rates, lower-income participants with lower attachment to their firms are more likely to forfeit unvested employer contributions. Our null retention result also stands in contrast to Cole and Taska (2023), who find that employer contributions to 401(k) plans are *more* important to firms' recruiting success than compensation paid as wages. The Cole and Taska (2023) finding is driven primarily by higher-income workers, who may place greater value on the tax advantages of 401(k) contributions than the lower-income workers who are most likely to forfeit 401(k) assets. Employer contribution benefits may also be more salient to job seekers than vesting schedules are to job separators.

The remainder of the paper is organized as follows. Section 2 provides institutional background on 401(k) vesting and Section 3 summarizes our administrative 401(k) data. Section 4 describes vesting-related forfeitures and their distributional effects. Section 5 tests for retention effects with the cross-plan difference-in-difference approach and Section 6 obtains within-plan retention estimates using our equivalency-method sample. Section 7 considers possible explanations for our null retention finding and presents results from the participant survey. Section 8 concludes.

2 Institutional background on 401(k) vesting

Contributions to 401(k) plans are of two broad types: employee contributions that participants defer from their paychecks, and employer (often matching) contributions that plan sponsors deposit directly into their participants' accounts. Participants always have full ownership of their own employee contributions. Plan sponsors, however, may enact vesting requirements that grant ownership of employer contributions only after specified time periods. Participants who separate from their employer before satisfying their plan's vesting requirements must forfeit some or all of the employer contributions they received during their time at the firm.

Vesting schedules specify the time frames over which participants earn ownership of employer contributions. These schedules are described by two parameters. First, the vesting period is the minimum employment tenure that participants must reach to earn full ownership rights. Vesting schedules can only be stated in annual terms, so the vesting period is always an integer number of years. Participants who remain employed through the end of the vesting period have full ownership of all employer contributions – those they received during the vesting period as well as ongoing contributions they receive after the end of the vesting period. Employers may choose not to enact vesting requirements; participants at these "immediate vesting" plans always have full ownership of employer contributions, no matter how long they have been employed.

The second vesting parameter is the type or shape of the schedule, which determines the rate at which participants accrue ownership rights during the vesting period. Under "cliff" vesting, all ownership rights are conferred at the end of the vesting period; under "graded" vesting, participants' ownership share increases gradually each year. For example, with three-year cliff vesting, participants employed for less than three years have a 0% ownership share in their employer contributions while participants employed for three years or more have a 100% ownership share. Under a three-year graded schedule, participants employed for less than one year have 0% ownership, those with tenure between one and two years have a nonzero ownership share (e.g., 33%), those with tenure between two and three years have a larger ownership share (e.g., 67%), and those with tenure of three years or more

have 100% ownership.[8] Under current federal regulations, the vesting period may be at most three years for cliff schedules and at most six years for graded schedules.

Vesting rules are enforced at the time of a participant's separation from his or her employer. At the time of separation, the 401(k) recordkeeper uses the plan's vesting schedule and the participant's employment tenure to determine the participant's ownership share in employer contributions. If the participant's ownership share is less than 100%, the recordkeeper executes a forfeiture transaction by withdrawing the unvested employer contributions (including all investment returns they have earned) and returning them to the employer.[9] The remaining market value of vested employer contributions, as well as the market value of the participant's own employee contributions, are disbursed to the participant (to be rolled over to an IRA or a new 401(k) plan, or taken as a cash distribution).

3 Data

3.1 Vanguard administrative data

We use administrative 401(k) data from Vanguard.[10] The database includes all information that Vanguard collects in the normal course of its recordkeeping responsibilities. At the plan level, we observe employer characteristics (such as industry) as well as plan rules (including vesting schedules). At the participant level, we observe hire and separation dates, vesting events (the dates on which a participant accrues an additional year of vested service), financial transactions (including contributions, withdrawals, and vesting-related forfeitures), account balances, and demographic information (including age).

Participant income is the only relevant variable that is not directly observable in the

[8]The exact path of graded ownership shares may differ between plans. For example, a three-year graded schedule may grant 33% ownership after one year and 67% ownership after two years; an alternative three-year graded schedule may grant 25% ownership after one year and 75% ownership after two years.

[9]The forfeitures that employers recover must be spent on plan-related expenses (including employer contributions for remaining participants).

[10]Our data also include other qualified defined-contribution retirement plans, such as 403(b) plans. 401(k) and 403(b) plans differ mainly in the legal status of the plan sponsor (for-profit vs. tax-exempt) and operate very similarly, so we do not need to treat them differently in our empirical analysis. For brevity, we refer throughout the paper to our sample as being composed of "401(k) plans."

administrative recordkeeping data. We do, however, have two ways to measure income. First, for about two thirds the plans it administers, Vanguard performs annual 401(k) non-discrimination testing mandated under ERISA. As part of the testing process, these employers provide Vanguard with annual compensation figures for their employees (those reported on IRS form W-2). Second, for a distinct but partially overlapping subset of about two thirds of plans, Vanguard tracks participants' elected contribution rates (the percentage of income that participants choose to defer as 401(k) contributions). Dividing participants' total 401(k) contributions in a given year by their average elected contribution rate in that year gives an estimate of their annual income. The directly reported income figures are more reliable than the income estimates we infer from contribution activity,[11] so we use directly reported income if it is available for a given participant-year and otherwise use our income estimate. Combining the two income sources in this way gives income information for about three quarters of the participant-year combinations in our sample; specifications that control for income limit to this subset.

We use the administrative data to form two distinct empirical samples. First, to study vesting-related forfeitures and their distributional effects, we form what we refer to as the "forfeiture sample." As discussed in Section 2, vesting-related forfeitures are only observable at the time of separation, when the recordkeeper splits the accumulated 401(k) balance among the participant and the employer according to the participant's vested ownership share. Our forfeiture sample thus conditions on separation and considers all participants who separated from their employer between 2010 and 2022 (the period during which the necessary components of the administrative data are available). We do not place any conditions on participants' hire dates, since conditioning on both hire and separation dates would require employment spells to begin and end within a fixed time period, which would mechanically decrease average employment tenures and mechanically increase

[11] Our inferred income estimates track W-2 figures closely for low- and middle-income employees but tend to give slight understatements for high-income employees (who are more likely to earn non-deferrable bonuses and hit annual limits on 401(k) contribution amounts). Since the source of income data is determined at the plan level (if a plan uses Vanguard for non-discrimination testing, it provides W-2 income data for *all* of its employees), the tendency for our inferred income estimates to understate W-2 income has a limited effect on within-plan income rankings (at plans without W-2 income, our estimates understate high earners' income to roughly the same extent, and high earners are still likely to be sorted into higher quintiles).

forfeiture rates. As a result, some participants in the forfeiture sample were hired before our administrative data become complete in 2010; we accept this because complete histories of employment spells are not strictly necessary to study forfeitures at separation.[12]

Second, to study the retention effects of vesting schedules, we form what we refer to as the "retention sample." When studying retention effects, it is important to observe the plan rules that prevailed at participants' hire dates, including vesting schedules and associated rules that govern participants' accrual of vested employment service. We also aim to avoid the censoring biases associated with estimating survival probabilities from incompletely observed employment spells (e.g., observing pre-2010 hires who remain employed long enough to appear in post-2010 data but missing pre-2010 hires who separate before 2010). We therefore restrict our retention sample to participants hired between 2010 and 2022.[13]

To remove composition effects and eliminate complications arising from plan entry and exit, we limit the retention and forfeiture samples to plans that are administered by Vanguard in both 2010 and 2022. This balanced panel includes 1,497 plans, with about 4.7 million separations during the 2010-2022 period. After imposing further sample restrictions (discussed in detail in Section 5), our cross-plan retention sample comprises about 1.45 million hires at 1,008 plans during the 2010-2022 period. Our within-plan retention sample, which restricts to 29 plans using the relatively uncommon equivalency method for service crediting (see Section 6.1), is substantially smaller, with 89,777 hires during the 2010-2022 period. Within the equivalency-method sample, 18,338 participants are hired on either side of the July-August service-crediting discontinuity. We give summary statistics for the forfeiture and retention samples in Table 1.

3.2 Plan-level vesting rules

Our administrative data show the vesting schedules in force at each plan as of the end of each calendar quarter. Vesting schedules change less frequently than other plan features,

[12]Even in cases where participants are hired before 2010, we can observe the hire date and therefore calculate tenure at separation.

[13]Our administrative data run through 2023, so we can observe at least one year of post-hire data for all participants in the retention sample.

Table 1: Summary statistics

	Forfeiture sample	Retention sample	
		Cross-plan (D-i-D)	Within-plan (Equiv. method)
Number of plans	1,497	1,008	29
Number of participants	4,762,709	1,451,656	89,777
Hired in July or August			18,338
Median age	36	33	30
Income			
25th percentile	26,501	27,740	46,769
Median	50,063	47,320	66,549
75th percentile	93,483	80,080	100,901
Participant-weighted vesting shares			
Immediate	35%	37%	0%
2-year cliff	5%	5%	5%
3-year cliff	15%	15%	5%
3-year graded	3%	1%	36%
5-year graded	15%	9%	24%
6-year graded	10%	7%	30%
Other	17%	26%	0%

Note: In the forfeiture sample, age, income, and vesting schedule are observed at the time of separation; in the retention sample, age, income, and vesting schedule are observed at the time of hire. If a given participant's plan has multiple vesting schedules, we associate the participant with the vesting schedule that has the longest vesting period. Income figures are given in real 2022 terms.

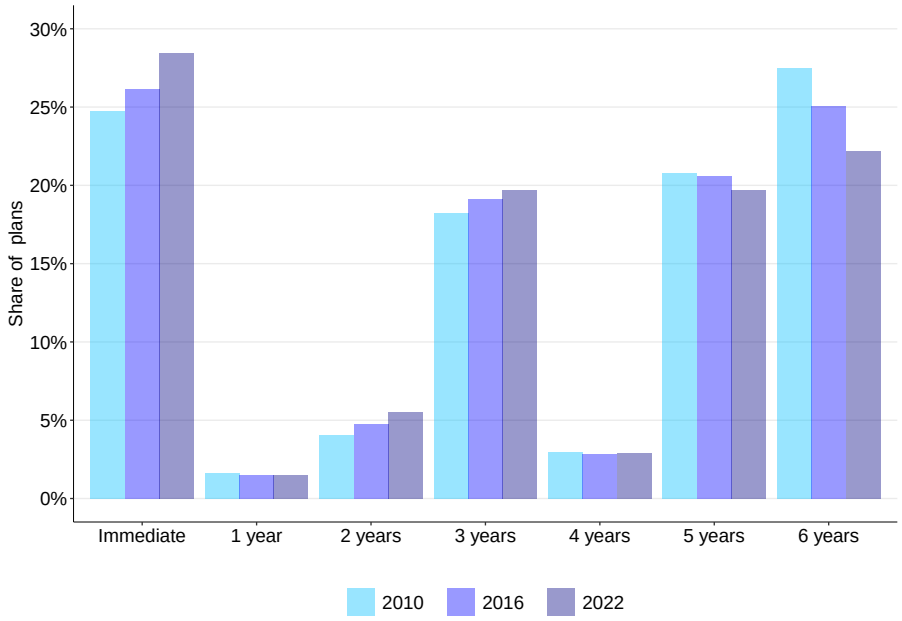
so the number of vesting rule changes within our balanced panel during the 2010-2022 period is quite small. Our retention analysis generally restricts to the (large majority) of participants whose plans did not make vesting changes during our sample period.

An important consideration in characterizing vesting rules is the fact that a given plan may have multiple vesting schedules in force at the same time. Plans may have multiple vesting schedules because plan sponsors may make more than one type of employer contribution and are free to apply a different vesting schedule to each contribution type. The most common such case is a plan sponsor that makes both a matching contribution and a non-matching contribution and chooses to use two different vesting schedules.[14]

When a given participant is subject to multiple vesting schedules and our retention analysis requires that we associate the participant with a single vesting schedule, we use

[14]Matching employer contributions are awarded as a function the participant's employee contributions. An example of a common matching contribution is one that awards $0.50 of employer contributions for each $1.00 of employee contributions up to 6% of the employee's income. Non-matching employer contributions do not depend on the employee's contribution activity. An example of a non-matching contribution is one that awards employer contributions equal to 3% of employees' income, to all employees at the firm (regardless of the employees' contribution rates and even for employees who do not contribute to the plan).

Figure 1: Distribution of vesting periods among the balanced panel of plans



Note: Vesting schedules are as of the end of the calendar year. For plans with multiple vesting schedules at the end of a given calendar year, we take the schedule with the longest vesting period.

the schedule with the longest vesting period. By doing so, we identify the minimum tenure the participant must reach to earn ownership of *all* employer contributions. We also note that this simplification is innocuous for most plans. In 2022, 59% of plans make a single type of employer contribution, 13% make two contributions but vest both immediately, and 11% make two contributions but subject both to the same non-immediate vesting schedule. Collapsing to the longest schedule thus makes no practical difference for 83% of plans in our sample. An additional 10% of plans make two contributions and subject the larger contribution (in terms of gross contribution flows) to a longer vesting schedule; for these plans, collapsing to the longest schedule ensures that we capture the vesting events with the most employer contribution dollars at stake. Finally, we show that our main difference-in-difference results in Section 5 are unchanged when we restrict to plans with a single type of employer contribution. The potential presence of multiple vesting schedules does not complicate our forfeiture analysis in Section 4: the entire set of employer contribution and vesting rules at a given plan is reflected in the forfeiture outcomes that we directly observe among separating participants.

Figure 1 shows the distribution of vesting periods among our balanced panel of plans

at the beginning (2010), middle (2016), and end (2022) of our empirical time frame. If a plan has multiple vesting schedules in a given year, we take the schedule with the longest vesting period. There is little change in the distribution over time, reflecting the fact that plans change their vesting rules infrequently. About 25% of plans use immediate vesting. Most of the remaining plans vest their participants over three years (the longest cliff schedule permitted under federal law), five years, or six years (the longest permissible graded schedule). Vesting periods of one, two, and four years are relatively uncommon, each accounting for less than 5% of plans in the sample.

Vesting schedules in our sample do not correlate strongly with observable firm characteristics. Figure A.1 shows firm industry prevalence by vesting schedule. Although industry composition differs slightly across vesting configurations, each schedule appears at a substantial number of firms across all major industries in our sample. We also calculate annual turnover rates (separations divided by actively employed participants) by vesting schedule. Figure A.2 shows that turnover rates hover around 10% during our sample period and do not appear to correlate systematically with employers' vesting choices.

We offer two potential explanations for the lack of observable firm differences across vesting schedules. First, firms that choose to obtain safe-harbor exemptions from nondiscrimination testing must generally vest their employees within two years. The decision to adopt safe-harbor standards (and the resulting constraints on vesting schedules) may be independent of industry and other firm characteristics. A second potential explanation is employer inertia. Firms introducing 401(k) plans were often guided by the design of their previous defined-benefit plans, and given the infrequency of vesting changes among 401(k) plans, it is likely that many of these initial vesting choices remain in force today.

4 The incidence of 401(k) forfeitures

4.1 Forfeitures among separating participants

Despite the prevalence of vesting requirements in 401(k) plans, relatively little is known about the vesting-related forfeitures that occur when participants separate from their

employers. What share of separations occur before the end of the vesting period and require a forfeiture of employer contributions? Are forfeitures more common among certain types of participants? How large are forfeitures relative to participants' total 401(k) balances? Our large sample of separations allows us to answer these questions.

As discussed in Section 3.1, the forfeiture sample includes separations that occur among our balanced panel of plans between 2010 and 2022. Although we define the sample based on participants' separation dates, forfeitures often occur with a lag after separation. For participants who separate with unvested balances, the forfeiture transaction occurs at the earliest of two events: when the participant's account is closed and their vested balances are withdrawn, or when five years have elapsed since separation.[15] Because our administrative data are complete through 2023, the potential five-year lag means that participants separating in 2018 are the most recent cohort for whom we can observe all forfeiture transactions. We present forfeiture statistics for the entire 2010-2022 period, but we take the 2018 figures as the most informative estimate of recent forfeiture rates.

Table 2 summarizes forfeitures for the first (2010) and most recent uncensored (2018) cohorts of separating participants, as well as for the overall 2010-2018 and 2010-2022 periods. Between 2010 and 2018, median tenure at separation among our balanced panel of plans falls from four years to two years. Falling employment tenures, along with secular growth in headcount, cause the number of separations to grow over time (more than doubling between 2010 and 2018). As shorter-tenured participants have become more likely to leave before the end of their vesting periods, forfeiture rates have increased. In 2010, 19.2% of separations involve a nonzero forfeiture of employer contributions; this share rises to 29.3% in 2018. Excluding plans with immediate vesting (where forfeiture rates are mechanically zero), the 2018 forfeiture rate is 40.2%.

Among participants with a forfeiture in 2018, the average forfeiture is $1,792 and accounts for 37.3% of the final account balance (a positive correlation between withdrawal

[15]The account closures that trigger forfeiture transactions can be voluntary or automatic (most plans automatically process immediate distributions for participants who separate with vested balances less than$5,000). The five-year waiting period for forfeiture transactions applies because workers who rejoin the firm within five years are generally entitled to restart their vesting service clock at their tenure as of the prior separation.

Table 2: Forfeitures among separating participants

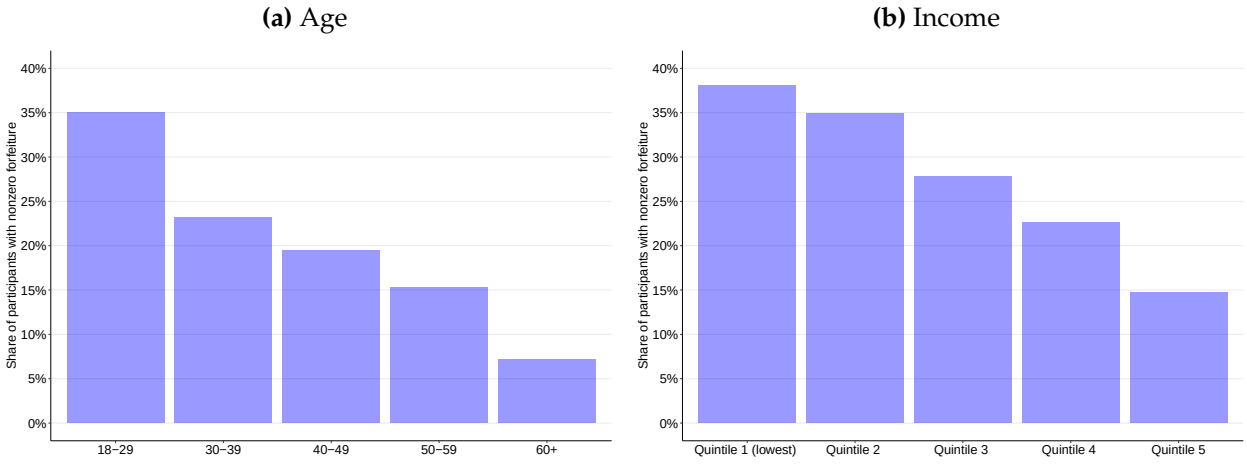
	2010	2018	Overall, 2010–2018	Overall, 2010–2022
Number of separations	202,589	444,838	2,763,219	4,762,709
Tenure at separation				
Mean	6.6	4.7	5.6	5.1
Median	4.0	2.0	3.0	2.0
Share with nonzero forfeiture	19.2%	29.3%	24.2%	23.6%
Among plans with vesting	29.4%	40.2%	35.6%	34.0%
Final balance				
With zero forfeiture				
Mean	86,987	83,214	90,985	83,434
Median	20,731	14,094	17,475	14,476
With nonzero forfeiture				
Mean	16,600	9,114	11,825	11,119
Median	4,175	1,721	2,384	2,198
Forfeiture share of balance (cond. nonzero)				
Mean	41.3%	37.3%	39.8%	37.8%
Median	33.0%	33.3%	33.4%	33.3%
Forfeiture share of income (cond. nonzero)				
Mean	4.8%	3.1%	3.6%	3.4%
Median	2.4%	1.4%	1.7%	1.6%
Forfeiture amount (cond. nonzero)				
Mean	3,623	1,792	2,487	2,148
Median	1,252	481	726	625
Sum (million)	141	234	1,666	2,419

Note: The final balance is the total market value of the participant's 401(k) account at the time of account closure. If account closure has not occurred, we use the market value of the participant's account at the time of separation. The forfeiture amount is the market value of employer contributions to which the participant does not have vested ownership rights and which are returned to the employer. Because forfeitures can occur with up to a five-year lag after separation, our data (which are complete through 2023) understate forfeiture rates for separations occurring in 2019 or later. Income is measured in the year of separation and annualized. Dollar values are given in real 2022 terms.

and forfeiture amounts causes the 37.3% figure to exceed the ratio of $1,792 and the mean account balance of $9,114). Expressed as a share of annual income rather than the final account balance, the mean forfeiture in 2018 is 3.1% of income and the median forfeiture is 1.4% of income. In Section 4.3, we use a lifetime simulation model to estimate the (significantly larger) cumulative effect of forfeitures on workers' age-65 retirement wealth.

Figure 2, which disaggregates the sample and plots forfeiture rates by age group and income quintile, demonstrates that forfeiture incidence is highly heterogeneous. Forfeiture rates are monotonically decreasing in both age and income, and the magnitudes of these relationships are large. Separating participants ages 18-29 have a forfeiture rate of 35%

Figure 2: Forfeiture rates by age and income



Note: The forfeiture rate is the share of separating participants who are not fully vested at separation and forfeit a nonzero amount of employer contributions.

during the 2010-2022 period, participants ages 30-59 show forfeiture rates in the 15%-25% range, and those above age 60 have forfeiture rates below 10%.[16] Similarly, when grouping separating participants into year-specific income quintiles, the forfeiture rate in the bottom quintile is 38%, versus 15% in the top quintile. Appendix Table A.1 formalizes these comparisons with simple linear regressions that control for age and income bins along with year and plan fixed effects; the age and income effects hold conditionally (i.e., do not arise solely from the correlation between age and income) and are robust to the inclusion of plan fixed effects.

The picture that emerges from Table 2 and Figure 2 is that vesting-related forfeitures are common occurrences, affecting about 30% of separating participants in the most recent cohort with complete forfeiture data. These forfeitures are concentrated in younger and lower-income participants, whose shorter tenures and lower incomes cause their account balances to be relatively low at the time of separation. As a result, the aggregate forfeiture amounts returned to employers are modest – among 2018 separators, aggregate forfeitures are $234 million, or about 3% of the gross employer contributions that plans with vesting requirements in our sample made in 2018. But forfeitures do represent meaningful dollar amounts for the affected participants, as they account for about 40% of these participants'

[16]This is partly because federal law requires participants to become fully vested at their plan's specified normal retirement age (usually 65).

final balances and forgo many years of tax-advantaged market returns.

4.2 Distributional effects

In this subsection, we quantify the distributional effects of vesting requirements and show that employer contributions are significantly more regressive when measured on a net basis (i.e., accounting for forfeitures that occur at separation) instead of a gross basis (i.e., not accounting for forfeitures). Gross employer contributions are likely more salient to participants. But since all participants eventually separate from their employers, net contributions – representing the dollars that participants can actually use to finance retirement consumption – are arguably more relevant.

To ensure that we can observe sufficiently long pre-separation histories of gross employer contributions, we limit to participants separating between 2016 and 2022 (the second half of the 2010-2022 sample period).[17] Additionally, because we aim to estimate distributional effects for treated participants who are actually subject to vesting requirements, we restrict to the roughly 75% of plans in our sample with non-immediate vesting. Finally, since non-participants do not appear in our recordkeeping data, differential participation rates by income do not affect our estimates – we describe the distribution of gross and net employer contributions conditional on plan participation.[18]

We sort separating participants into plan-year-specific income quintiles based on their annualized income in the year of separation. Because we aim to rank separating participants relative to the full income distribution at the plan (and not just relative to other separating participants), we use all active plan participants to define the cut points for the income quintiles. We then look back at separating participants' entire observable history of income and employer contributions. Since our administrative data stretch back to 2010, this

[17] Recall that the longest vesting schedule permitted under federal law is six years. By considering separations from 2016 onward, we ensure that the entire employment spell for any participant experiencing a forfeiture (who must have been hired in 2010 or later) is observable. This in turn ensures that the annualized net employer contribution rates we compute in (2) are accurate. For separations before 2016, (2) may allocate a nonzero forfeiture f_i across only a subset of the years the participant was employed at the plan.

[18] The results are also robust to dropping the 2019-2022 separators whose forfeitures are partially time-censored (implying that the share of forfeitures that are censored does not differ too much across income groups).

observable history begins either in 2010 or in the year the participant was hired at the plan, whichever is later. We use the income and contribution history to define:

$$G_i := \frac{\sum_{t=t_{i_0}}^{T_i} c_{it}}{\sum_{t=t_{i_0}}^{T_i} y_{it}}, \tag{1}$$

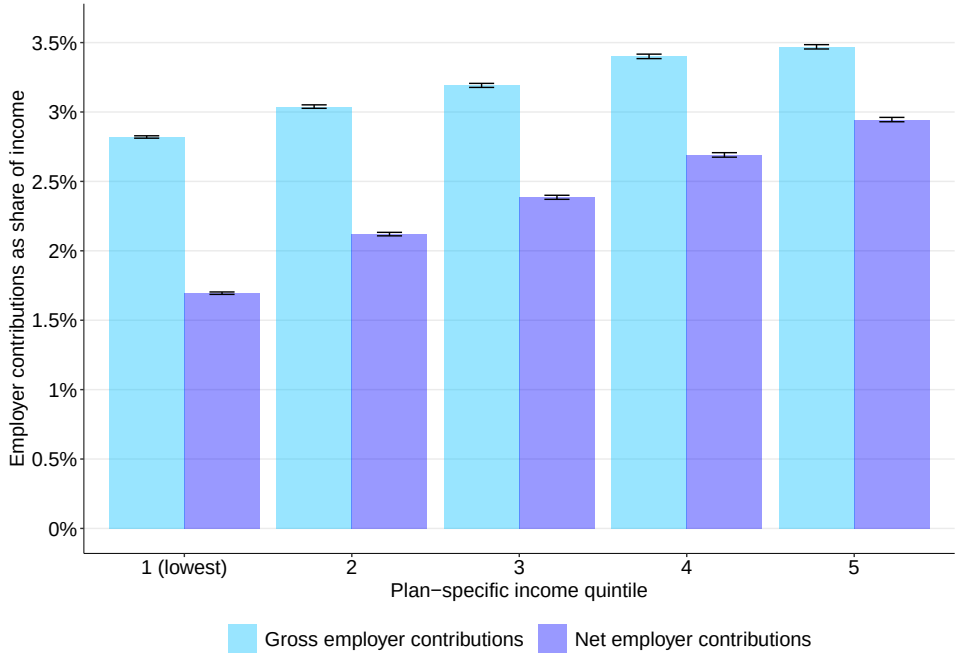
$$N_i := \frac{(\sum_{t=t_{i_0}}^{T_i} c_{it} - f_{it}^{pv})}{\sum_{t=t_{i_0}}^{T_i} y_{it}}, \tag{2}$$

where i indexes participants, t_{i_0} is the first year of participant i's observable history, T_i is participant i's separation year, c_{it} is the gross amount of employer contributions awarded to participant i in year t, and y_{it} is participant i's income in year t. The final term f_{it}^{pv} is the amount of employer contributions that participant i forfeits at separation (which may be zero), converted back to year-t present-value terms.[19] G_i thus gives the average annual gross employer contributions that participant i receives during the employment spell, scaled by his or her average annual income. N_i also gives average annual employer contributions as a share of income, but is a net measurement that adjusts for any balances forfeited at separation. We estimate the distributional effects of vesting requirements by comparing the distributions of G_i and N_i across plan-specific income quintiles.

Let $\overline{G}_q$ and $\overline{N}_q$ be the average values of G_i and N_i for participants in income quintile q. Figure 3 plots the values of $\overline{G}_q$ and $\overline{N}_q$, along with their 95% confidence intervals. The expected pattern clearly emerges, with net employer contributions tilted more heavily toward high-income participants than gross contributions. Because higher-income participants are more likely to maximize employer matches, the gross distribution is itself regressive, rising from 2.82% in the bottom quintile to 3.47% in the top quintile. Net contributions, however, are even more sharply increasing in income. The high rate of forfeitures among the bottom quintile brings average annual net contributions to 1.69% of income, a decrease

[19] To make this present-value adjustment, let $s_{it} = \frac{y_{it}}{\sum_t y_{it}}$ be the share of participant i's total observable income earned in year t. We allocate the total forfeiture amount f_i into annual components f_{it} in proportion to the annual income shares, so $f_{it} = s_{it} f_i$. We then calculate present-value-adjusted annual forfeiture amounts f_{it}^{pv} by dividing by a constant discount factor: $f_{it}^{pv} = \frac{f_{it}}{(1+r)^{(T_i - t)}}$, where T_i is the participant's year of separation and r is a discount rate approximating average annual returns (we set $r = 0.05$).

Figure 3: Comparison of gross and net employer contributions by plan-specific income quintile



Note: This graph considers participants who separate during the 2016-2022 period. Income quintiles are plan-year-specific and are defined using all active plan participants. Separating participants are sorted into income quintiles based on their annualized income in the year of separation. Gross and net employer contributions are as defined in (1)-(2). We restrict to plans with vesting requirements (about 75% of the plans in our sample). Whiskers give 95% confidence intervals.

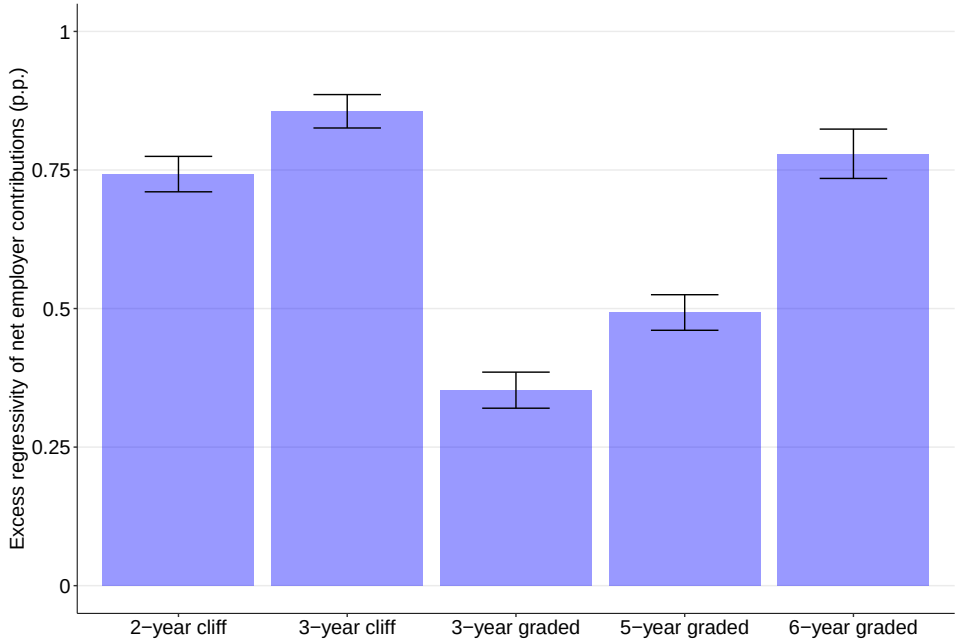
of 1.13 percentage points from the corresponding gross figure. The gap between gross and net contributions shrinks in each successive quintile, with the highest-income participants showing only a 0.52-percentage-point gap between the gross and net figures. When comparing the top and bottom quintiles, the difference is 0.65 percentage points in gross terms but 1.25 percentage points in net terms.

Formalizing the comparison between the top and bottom income quintiles, we define:

$$R := (\overline{N}_5 - \overline{N}_1) - (\overline{G}_5 - \overline{G}_1). \tag{3}$$

The first term in (3) is the difference between the top and bottom quintiles in average annual net contributions, and the second term is the top-bottom difference in gross terms. R, which we refer to as the "excess regressivity" of net employer contributions, is a summary measure describing the distributional effect of vesting requirements. The global averages shown in Figure 3 imply a sample-wide excess regressivity value of 0.6 percentage points.

Figure 4: The excess regressivity of net employer contributions, by vesting schedule



Note: This graph considers participants who separate during the 2016-2022 period. The excess regressivity of net employer contributions is as defined in (3). Results are shown separately for five common vesting schedules in our sample. If a separating participant's plan makes multiple types of employer contributions and has multiple vesting schedules, we associate the participant with the longest vesting schedule at the plan. Whiskers give 95% confidence intervals.

Figure 4 plots the excess regressivity measure R separately for five common vesting schedules in our sample. Recall that we associate a participant with the longest vesting schedule at his or her plan if the plan has multiple schedules. The heterogeneity by schedule is intuitive, with cliff schedules (which confer all ownership rights at the end of the vesting period) showing stronger distributional effects than graded schedules (which allow participants separating before the end of the vesting period to retain partial ownership of employer contributions). Within the cliff and graded categories, excess regressivity also increases with the length of the vesting period.

A potential concern with the preceding analysis is our classification of separating participants into plan-specific income quintiles. If longer employment tenures themselves tend to raise income (through skill growth, outside-offer bargaining, or other mechanisms), the relationships documented in Figures 3-4 could be partly mechanical: in this case, long employment tenures would tend to push participants into higher income quintiles in the year of their separation and simultaneously reduce their forfeiture rates. Ideally, we would classify separating participants based on a measure of their lifetime income, but

we are constrained in only observing income earned during their employment at the plan. Even so, we present robustness results demonstrating that our estimates of distributional effects are unlikely to be mechanical. Appendix Figure A.3 reproduces Figure 3 but sorts separating participants into quintiles based on their average annual income at the plan rather than their final annual income. Figure A.4 shows a second robustness test where we classify separating participants based on their first observable annual income at the plan. These alternative approaches weaken any underlying link between tenure and income but produce results that are qualitatively similar to the baseline analysis in Figure 3. Finally, Figure A.5 shows that the results are robust to dropping the 2019-2022 separators whose forfeiture transactions are partially time-censored.

4.3 Lifetime costs

For most workers in our recordkeeping data, we can observe only a single separation and forfeiture outcome. The lifetime costs of forfeitures are likely much larger than the preceding estimates, for two reasons. First, workers hold multiple jobs during their careers and are at risk of a forfeiture each time they separate. Second, the cost of any given forfeiture in terms of potential retirement consumption grows over time as the forfeited funds forgo compounded market returns. These forces are particularly important in our setting, with secular decreases in employment tenure raising the average worker's job-transition rate (see Table 2) and forfeitures occurring most often among young workers with long investing horizons (see Figure 2). To assess the lifetime costs of forfeitures, we conduct a simple simulation with inputs calibrated to important features of the administrative data.

4.3.1 Simulation process

We model the career of a worker during the prime working ages of 25 to 59, with retirement wealth earning market returns until age 65. To evaluate the compounded future value of forfeiture losses, we compare the worker's retirement wealth at age 65 to a counterfactual where any forfeitures the worker experienced during her career did not occur (but where all other realized simulation outcomes are the same).

Mirroring the setup in Section 4.2, we generate separate simulation results for each income quintile. Using age-dependent inputs that are specific to each income quintile, we conduct 10,000 lifetime simulations for each quintile and report the distribution of forfeiture outcomes across the simulations.

Appendix B provides a detailed discussion of the inputs to the simulation model (separation hazards, income paths, participation and contribution rates, and the distribution of employer contribution and vesting rules), which we estimate directly from the administrative recordkeeping data. The key inputs are the income-, age-, and tenure-specific separation hazard rates. Because lower-income workers separate at higher rates, they hold more jobs during their careers, have shorter tenures at separation, and are more likely to suffer forfeitures.

4.3.2 Results

Table 3 summarizes the simulation results. Lower-income workers' higher turnover rates cause them to experience more forfeitures over the course of their careers: in the bottom quintile, 50% of simulated workers have at least two forfeitures, 25% have at least three, and 10% have at least four. In contrast, the median number of forfeitures among the top quintile is zero and fewer than 10% of top-quintile workers experience two or more.

Accounting for foregone returns makes the lifetime costs of vesting requirements significantly larger than raw forfeiture amounts. The age-65 values in the bottom panel of Table 3 are generally about seven times larger than the corresponding forfeiture amounts in the middle panel, which implies about 30 years of compounding (i.e., from age 35 to 65) given average annual real TDF returns of about 7%. Among the bottom income quintile, the median value of forfeiture losses at age 65 is $16,256 and the average value is $23,229. These dollar amounts are substantial particularly when considered relative to earnings: the average forfeiture loss represents 62% of the bottom quintile's lifetime average annual income. Higher-income workers are less likely to experience forfeitures and total losses are smaller relative to their earnings, but because employer contributions scale with income, losses among high earners can be significant in dollar terms when they do occur. For

Table 3: Lifetime simulation results

	Income quintile				
	Bottom	Second	Third	Fourth	Top
Number of forfeitures					
Percentiles					
10th	0	0	0	0	0
25th	1	0	0	0	0
50th	2	1	1	0	0
75th	3	2	1	1	1
90th	4	2	2	2	1
Mean	1.86	0.98	0.82	0.76	0.44
Sum of forfeitures					
Percentiles					
10th	0	0	0	0	0
25th	577	0	0	0	0
50th	2,431	1,221	467	0	0
75th	4,834	4,438	4,838	6,806	6,755
90th	7,576	8,243	9,341	13,475	18,004
Mean	3,249	2,896	3,073	4,298	5,023
Age-65 value of forfeitures					
Percentiles					
10th	0	0	0	0	0
25th	3,045	0	0	0	0
50th	16,256	6,901	1,407	0	0
75th	34,603	36,728	35,302	44,179	30,039
90th	56,516	68,124	76,098	101,085	108,732
Mean	23,229	23,457	24,092	30,446	29,030
As share of avg. annual income	62%	39%	29%	26%	15%

Note: We conduct 10,000 simulations of the model for each income quintile and report the distribution of forfeiture outcomes across the simulations. See Section 4.3 and Appendix B for a description of the simulation model and its parameters.

example, the median top-quintile worker does not experience a single forfeiture but losses are $30,039 at the 75th percentile and $108,732 at the 90th percentile. Taking a global view, the average age-65 value of forfeiture losses across all quintiles is $26,051, which is 26% of cross-quintile lifetime average annual income.

The regressive forfeiture costs documented above are only one empirical consequence of vesting requirements. In the remainder of the paper, we consider a second potential effect that may deliver benefits for firms and their workforces: employee retention.

5 Estimating retention effects with cross-plan variation

5.1 Identification strategy

Our first identification strategy to estimate the retention effect of vesting schedules is a difference-in-difference approach. A simple retention test would be to determine whether voluntary separation probabilities jump after discrete vesting events. But directly attributing any bunching of separations after vesting deadlines to the effect of 401(k) vesting schedules would not be appropriate, as there are likely to be confounding treatments occurring around the same dates.

Most 401(k) plans calculate vested employment service as the number of years elapsed since the participant's hire date or plan-entry date. We refer to the date at which plans start counting elapsed time (hire or plan entry) as the "start date" for the rest of this section. As a consequence of elapsed-time service crediting, vesting dates generally coincide with participants' annual employment anniversaries. Simple before-and-after comparisons or bunching estimators would therefore be biased by the effects of non-vesting treatments that tend to occur at yearly tenure intervals. Possible confounds include wage increases and promotions; workers who do not receive expected raises or promotions might be more likely to quit (ADP Research, 2023). Another important confound is deferred compensation awarded outside of the 401(k) plan, including bonuses, company stock, and options. These forms of non-401(k) deferred compensation, which are prevalent among managerial and higher-income workers, could be subject to their own vesting schedules. If these vesting schedules coincide with the vesting schedules for 401(k) compensation, we might be assigning the effects of the former to the latter. Since we cannot observe vesting rules for non-401(k) compensation, we cannot control directly for any overlap.

A difference-in-difference design allows us to account for these potential confounds, as long as the usual identifying assumptions hold. We use participants in plans with immediate vesting as a natural control group.[20] The identifying assumptions for difference-

[20] An alternative approach would be a staggered-adoption specification where the source of variation is differential vesting durations among plans with vesting (e.g, using participants approaching their third employment anniversary at two-year cliff plans as a control group for participants approaching the end of a

in-difference designs are i) parallel trends and ii) no anticipation of treatment. The parallel-trends assumption posits that absent the treatment, the treated and control groups would have followed similar trends in the outcome variable. The conceptually distinct no-anticipation assumption requires that the treatment have no effect in periods prior to implementation.

The no-anticipation assumption merits careful consideration in our setting. Sophisticated participants with foresight of vesting dates may delay separations that otherwise would have occurred in pre-treatment periods until after the vesting deadline. But in this case, the delayed separations occurring just after the vesting deadline would manifest as positive difference-in-difference estimates during post-treatment periods and would lead us to reject the null hypothesis of zero retention effects. In other words, violation of the no-anticipation assumption in our setting is really a treatment effect that would tend to amplify separation spikes at vesting relative to non-vesting plans. In any event, we show below that there is no visual evidence of pre-trends or anticipatory behavior in the year before treatment. The lack of observable pre-trends gives us confidence in our research design and in our conclusion of zero retention effects.

Starting from our balanced panel of plans, we restrict to plans that use elapsed-time service crediting. This filter nets 1,008 plans with 1,451,656 participants. For the main analysis in this section, we focus on estimating the treatment effect for participants with three-year cliff vesting. The three-year cliff schedule is one of the most common vesting configurations and gives participants the strongest incentives to wait to vest before quitting. Compared to other cliff schedules, the incentives are larger because the accumulated balance is larger (cliff schedules may be at most three years). Compared to graded schedules, the incentives are larger because the jump in ownership share is larger (all jumps in graded schedules are strictly smaller than the 0-to-100% jump in cliff schedules). Though we focus on the three-year cliff, we also estimate treatment effects for other vesting schedules and obtain similar results.

three-year cliff schedule). We prefer to use never-treated participants (i.e., participants at non-vesting plans) as the control group to avoid several econometric difficulties that the difference-in-difference literature has identified regarding staggered-adoption designs (Baker et al., 2022).

For the main analysis with three-year cliff vesting, we use 181 plans in the treatment group (162,788 participants) and 428 plans in the control group (439,309 participants) for which we have income information. As shown in Appendix Table A.2, the two groups are roughly balanced across observable covariates, although the treatment group appears to have slightly higher incomes. The income difference is likely due to industry composition, as the treatment group contains more participants in higher-income industries (e.g., the Business/Professional and Finance classifications). We control for income and other observable covariates in our specifications.

5.2 Results

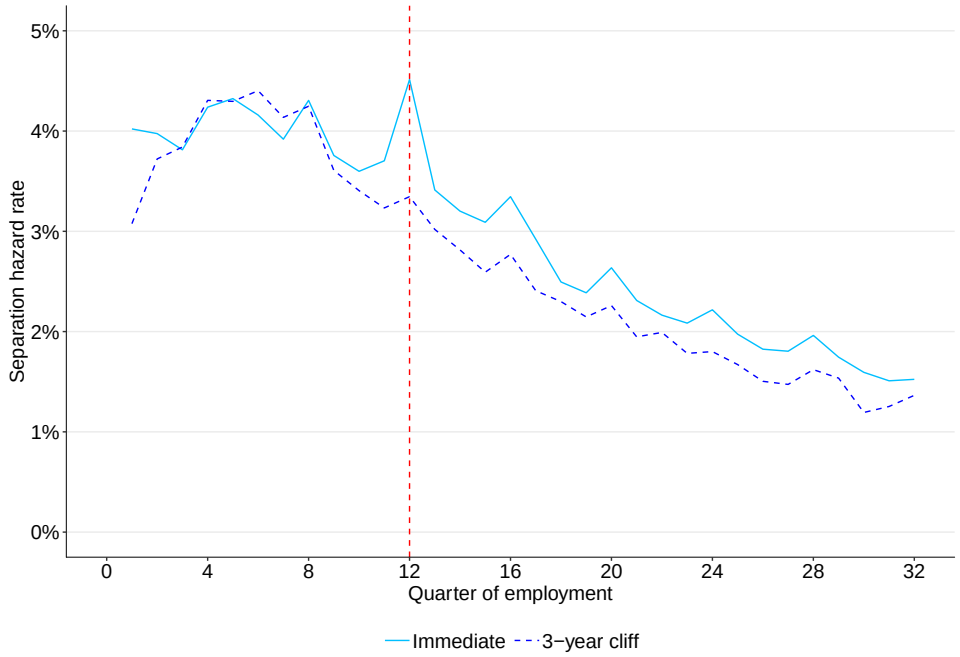
We begin by estimating voluntary separation hazard rates nonparametrically at a quarterly frequency. The hazard rate in a given quarter q is the probability of voluntarily separating in quarter q, conditional on having survived through the previous quarter $q-1$. We focus on voluntary separations because they are most relevant in testing the retention hypothesis: retention incentives aim to deter worker-initiated quits. The recordkeeping data contain a field explicitly identifying layoffs and firings, but plan sponsors are not required to flag involuntary separations to Vanguard.[21] It is therefore possible that our main sample contains some involuntary separations that were not flagged as such. We do not expect unflagged involuntary separations in the neighborhood of the third employment anniversary to be common enough to meaningfully affect our results, and for completeness we show in Appendix Figure A.9 that the null event-study result also holds for the identifiable involuntary separations that we drop from the main analysis.[22]

Figure 5 shows the voluntary separation hazard functions for the treatment group

[21] Separation reasons are not essential recordkeeping information because even involuntary separations during the vesting period incur forfeitures. About 5.5% of separations during the 2010-2022 period are explicitly coded as involuntary.

[22] If firms laid off workers just before vesting deadlines in order to recoup forfeitures, then our event study would yield positive pre-trends (unlike the flat pre-trends that we find in Figures 6 and A.9). Alternatively, if layoffs occur because of firm-wide or macroeconomic conditions that are unrelated to individual employees' vesting statuses, then our focus on voluntary separations is conservative (dropping separations that we know to be unrelated to vesting deadlines makes it harder to establish our null retention finding).

Figure 5: Voluntary separation hazard curves for the treatment and control groups



Note: The hazard rate is the probability of separating from the firm in a given quarter, conditional on having survived through the previous quarter.

(three-year cliff vesting) and control group (immediate vesting). The vertical dashed line marks the 12th quarter since the start date, at which point the treatment group vests. The hazard curve for the treatment group shows a noticeable spike at quarter 12, which might in isolation suggest a vesting retention effect. However, we also see a pronounced spike at the same quarter for the control group that does not face 401(k) vesting incentives.[23] The control spike likely captures the confounding treatments discussed above, including promotions, pay raises, and non-401(k) deferred compensation. Another important result in Figure 5 is that the spike at quarter 12 is not unique; spikes occur for both groups at regular annual intervals from the start date (quarters 8, 12, 16, ...). The annual spikes indicate that workers appear to separate with regularity at yearly intervals that are not systematically related to 401(k) vesting. This feature of the quitting hazard curve makes clear the importance of using an appropriate control group in order to isolate the effect of 401(k) vesting.

To test for retention effects, we estimate the following linear difference-in-difference

[23]The control spike is actually larger in raw terms in Figure 5; when we control for observables in the formal event-study analysis in Figure 6, the differential spike disappears.

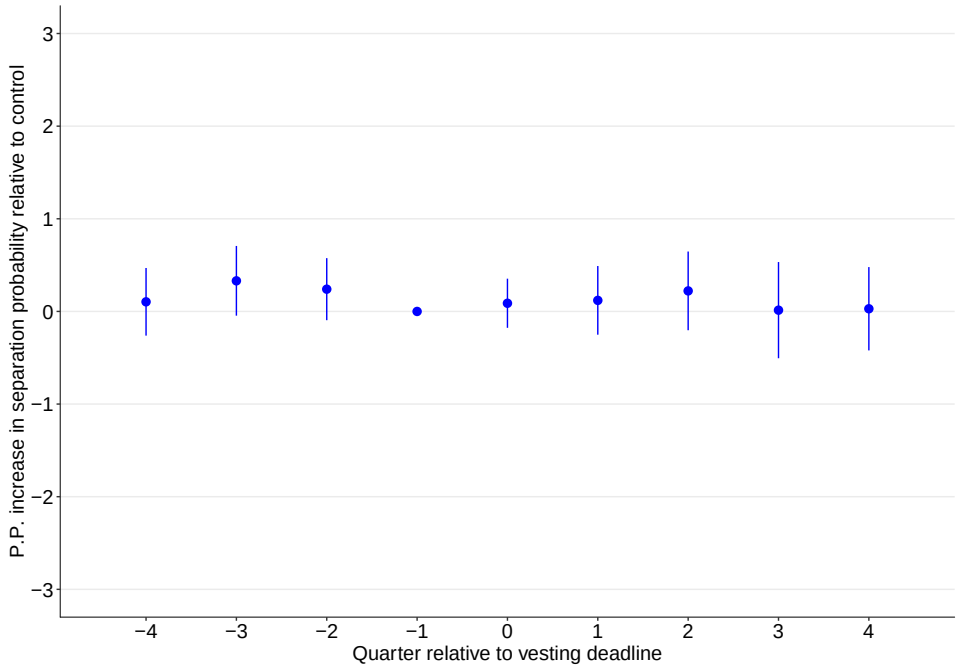
event-study specification, using four quarters before and after the vesting event:

$$\text{Quit}_{iq} = \sum_{k=-4}^{k=4} \beta_k \text{I}(t_q = \text{VQ} + k) * T_i + \sum_{k=-4}^{k=4} \theta_k \text{I}(t_q = \text{VQ} + k) + \alpha T_i + \gamma X_i + \varepsilon_{iq}. \quad (4)$$

Time t_q is indexed by quarters. VQ is the quarter in which the vesting event occurs, in this case quarter 12. The covariate vector X_i includes an age quadratic, log income, calendar-year and quarter fixed effects, industry fixed effects, and plan-size-decile fixed effects. T_i is an indicator equal to one for participants in treatment (three-year cliff) plans. The parameters of interest are β_k, which give the differences in separation hazard rates between treatment and control in each relative quarter k. We cluster standard errors at the plan level. We use a linear specification because it relaxes the computational constraints involved in estimating a nonlinear specification over a large sample of participant-quarter observations. For robustness, we also estimate a nonlinear complementary log-log separation model for the three-year cliff sample and similarly obtain null results (see Figure A.6). The results also hold when we restrict to plans that make a single type of employer contribution and thus have only a single vesting schedule (see Figure A.7, Table A.3 and the discussion in Section 3.2).

Figure 6 shows our main event-study results. All coefficients are normalized with respect to the $k = -1$ coefficient. First, we find no evidence of differential pre-trends before the treatment group vests. Workers in the treatment and control groups separate at roughly the same rates before the vesting event, which supports the parallel-trends assumption necessary for identification in difference-in-difference studies. Furthermore, we do not see any evidence of anticipatory behavior, since the coefficient at $k = -1$ is roughly the same magnitude as the rest of the pre-treatment coefficients. Second, we find no evidence of differential changes in voluntary separation probability after vesting: the post-treatment estimates are flat and centered at zero. Pooling all relative quarters $k \geq 0$, the combined effect of becoming vested in 401(k) employer contributions is -0.076 percentage points (with a standard error of 0.123 percentage points). This effect is statistically and economically insignificant; with a baseline pre-treatment quarterly quit probability of 3.5 percentage

Figure 6: Difference-in-difference event study for 3-year cliff vesting schedules

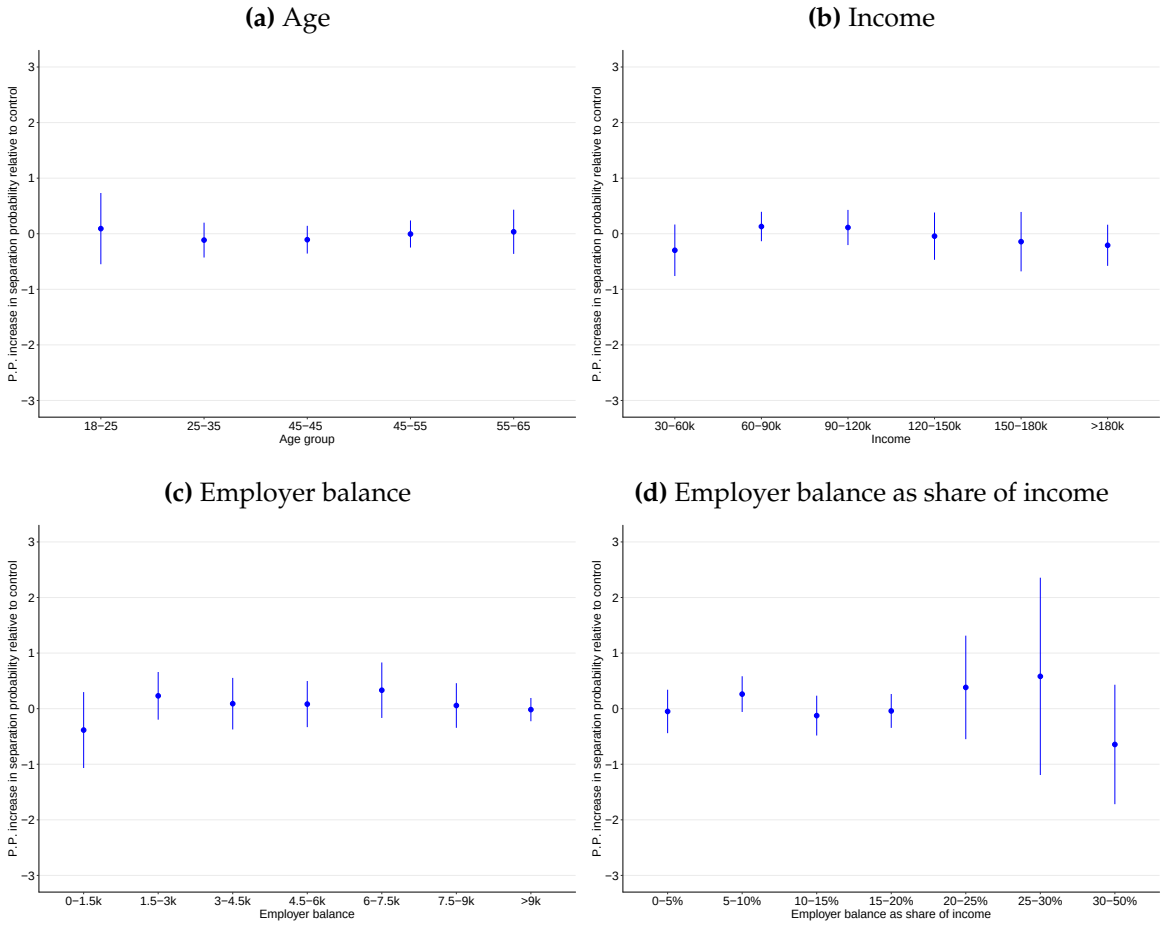


Note: We estimate a difference-in-difference event study with participants in three-year cliff vesting plans as the treatment group and participants in immediate-vesting plans as the control group. We plot the coefficient estimates β_k from (4) along with their 95% confidence intervals. The coefficient estimates are normalized with respect to relative quarter $k = -1$. The outcome variable is an indicator equal to one if the participant voluntarily separates in a given quarter. We control for an age quadratic, log income, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are clustered at the plan level.

points, the null estimate represents a decrease of 2.1% in the baseline quit hazard.

Next, we examine heterogeneity in treatment effects across age and income groups. We estimate a combined pre-post treatment effect for each demographic group using the same controls as in the main specification (4). Figure 7 shows that the null result holds across the age distribution. The effect is also consistently zero across the income distribution, including for more affluent participants who may have larger employer contribution balances at stake around the vesting event. We test the effect of balances directly in Figures 7c-7d. We measure the total accumulated balance in employer contributions at relative quarter $k = -4$ and split the sample according to the size of this balance, both in absolute terms and as a share of annual income. We expect the effects of vesting incentives, if they exist, to be larger for participants with larger employer contribution balances. But again, we find that the null effect holds across the distributions of absolute and relative employer balances. While the mean forfeiture amount during the 2010-2022 period is about $2,100,

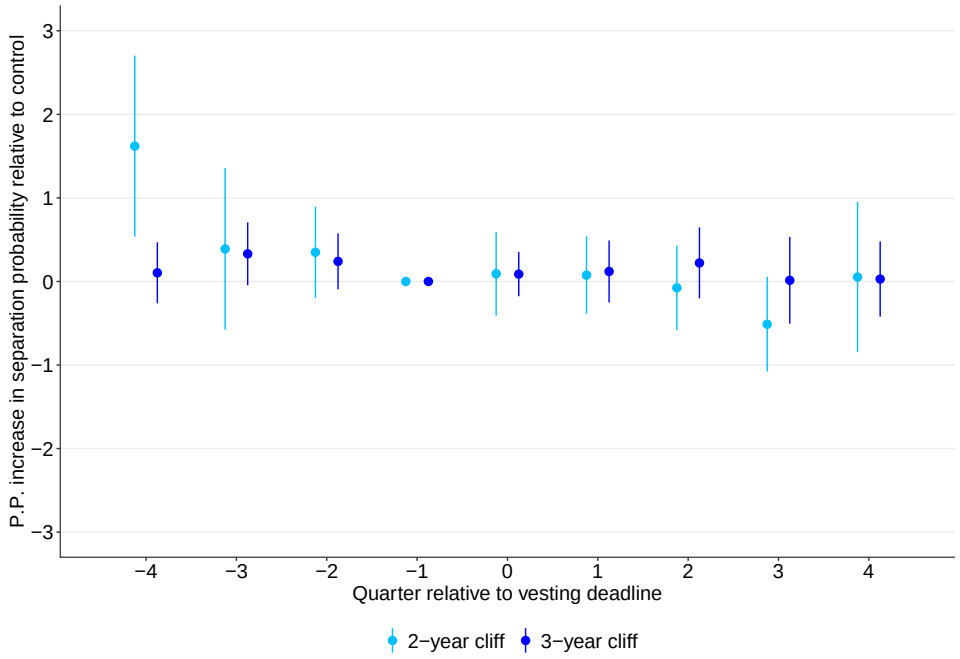
Figure 7: Heterogeneity in difference-in-difference estimates



Note: We estimate a difference-in-difference treatment effect with participants in three-year cliff vesting plans as the treatment group and participants in immediate-vesting plans as the control group. We estimate the treatment effect separately for each age, income, and employer contribution group. We measure employer contribution balances at relative quarter $k = -4$, in absolute terms and as a share of annual income The outcome variable is an indicator equal to one if the participant voluntarily separates in a given quarter. We control for an age quadratic, log income, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are clustered at the plan level.

the null effect is present even for participants who stand to lose double or triple that amount. The null effect is also present for participants who stand to lose employer contributions worth more than 10% of their annual income. This result is even more surprising when considering that the median wage gain for job switchers held between 4% and 5% during the 2010-2022 period (Federal Reserve Bank of Atlanta, 2025). Workers thus do not appear to react to vesting incentives even when they stand to lose 401(k) compensation amounts similar in size to the wage gains they might expect to realize from switching jobs (see Section 7.2 for a detailed analysis of wage gains among job swtichers).

Figure 8: Difference-in-difference event study for cliff vesting schedules

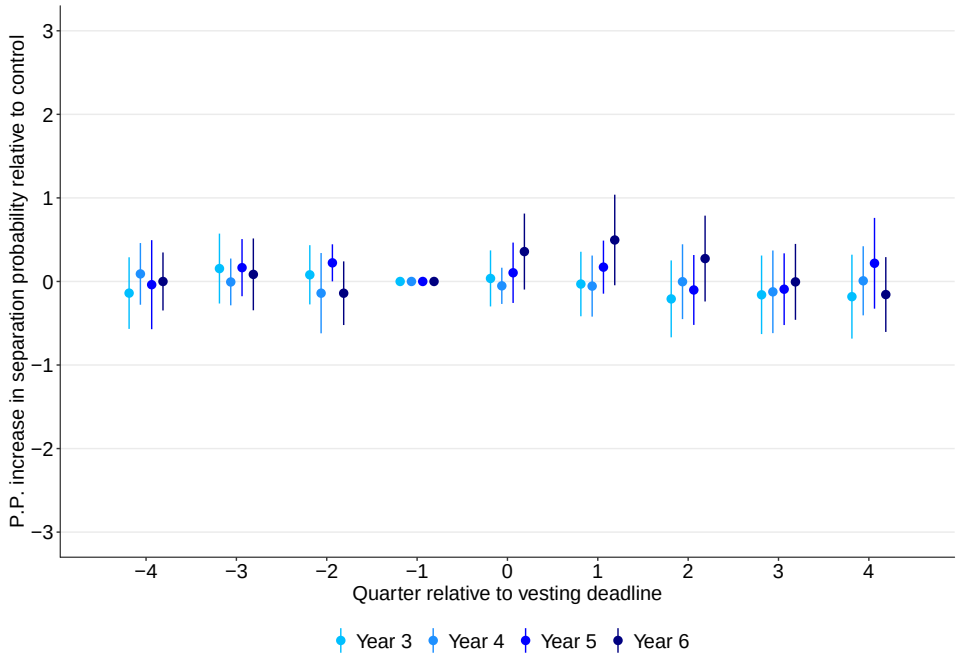


Note: We estimate difference-in-difference event studies with participants in three-year cliff and two-year cliff vesting plans as treatment groups and with participants at immediate-vesting plans as the control group. The outcome variable is an indicator equal to one if the participant voluntarily separates in a given quarter. We control for an age quadratic, log income, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are clustered at the plan level.

Finally, we estimate treatment effects for vesting schedules other than the three-year cliff. In Figure 8, we compare three-year cliff plans to two-year cliff plans. The two-year cliff estimates show some evidence of pre-trends at relative quarter $k = -4$ but are similar to the baseline $k = -1$ hazard for the rest of the pre-treatment period. Although the standard errors are larger because of the smaller sample size, we again find no evidence of increased separation probabilities following vesting. Given that monetary incentives at the two-year mark are likely smaller than those at the three-year mark, this result is not surprising.

We also estimate treatment effects for graded schedules. Estimating our main specification for graded schedules is more difficult because it is not obvious which vesting deadlines we should consider. For example, the six-year graded schedule, one of the most common graded durations, has vesting events not only at year six (when participants reach 100% ownership share) but also at every tenure anniversary starting in year two. We could focus solely on the treatment effect at the last vesting event, but this approach would not capture potential retention effects at other points during participants' employment spells.

Figure 9: Difference-in-difference event study for graded vesting schedules



Note: We estimate difference-in-difference event studies at yearly intervals (three, four, five, and six years) from the start date. We stack these event studies across each graded schedule and estimate a uniform set of interaction coefficients β_k. The outcome variable is an indicator equal to one if the worker voluntarily separates in a given quarter. We control for an age quadratic, log income, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are clustered at the plan level.

Our approach is to combine vesting deadlines in graded schedules that occur at the same point in participants' employment tenures. For example, for the three-, four- , five-, and six-year graded schedules, we can construct an event study that pools all participants under these schedules and is centered around the vesting event at quarter 12. For the three-year graded schedule, the three-year mark represents the last vesting event in the schedule; for participants at six-year graded plans, vesting events remain in the fourth, fifth, and sixth years. To estimate this event-study specification, we stack observations using data from the different graded schedules, following similar stacking approaches in the difference-in-difference literature (Cengiz et al., 2019; Baker et al., 2022). Since hazard rates decrease over the course of employment, it is important to combine data from the same portions of job spells to ensure the compatibility of the estimated treatment effects. This approach also ensures that employer contribution balances – and thus the stakes of the separation decision – are similar, since at each annual vesting event we are considering participants who have accumulated 401(k) balances for the same amount of time.

We identify the set of participants who are still employed four quarters prior to each annual vesting event (i.e., participants' third, fourth, fifth, and sixth employment anniversaries). Then, for each annual vesting event, we create a stacked event-study dataset [24] that combines data for the different graded schedules and estimate the following specification:

$$\text{Quit}_{iqs} = \sum_{k=-4}^{k=4} \beta_k \text{I}(t_q = \text{VQ} + k) * T_i + \sum_{k=-4}^{k=4} \theta_{ks} \text{I}(t_{qs} = \text{VQ} + k) + \alpha_s T_{is} + \gamma X_i + \varepsilon_{ijq} \quad (5)$$

We use the same set of control variables as in the main specification in (4) but now have intercept shifts α_s and baseline hazard rates θ_{ks} indexed by graded schedule s. The interaction effects β_k are the parameters of interest and are uniform across schedules.

Figure 9 shows the event-study results for graded schedules. Similar to our main results for cliff schedules, we find no differential voluntary separation trends before or after vesting events at years three, four, or five. At year six, there is a slightly positive but statistically insignificant point estimate (see Table A.4). These stacked results yield no evidence of economically or statistically significant retention effects at any annual vesting event within the graded schedules.

Taking into account all of the difference-in-difference results, our cross-plan analysis indicates that 401(k) vesting incentives do not exert causal retention effects. The null result holds across demographic groups, for different potential forfeiture amounts, and for different vesting schedules. In the next section, we corroborate this null result using an alternative research design based on within-plan variation.

6 Estimating retention effects using within-plan variation

6.1 Identification strategy

While most 401(k) plans (such as those studied in Section 5) use elapsed time since hire to determine participants' vested service, a minority of plans credit service using

[24] The stacked event-study dataset includes separate control and treatment groups for each graded vesting schedule, allowing for schedule-specific baseline hazards while estimating uniform dynamic treatment effects across schedules.

the "equivalency method."[25] Under the equivalency method, participants who work at least one hour during a specified calendar period (e.g., a month) are credited with a fixed number of service hours. If a participant is credited with at least 1,000 such service hours in a given calendar year, then that calendar year counts as a year of vested service.

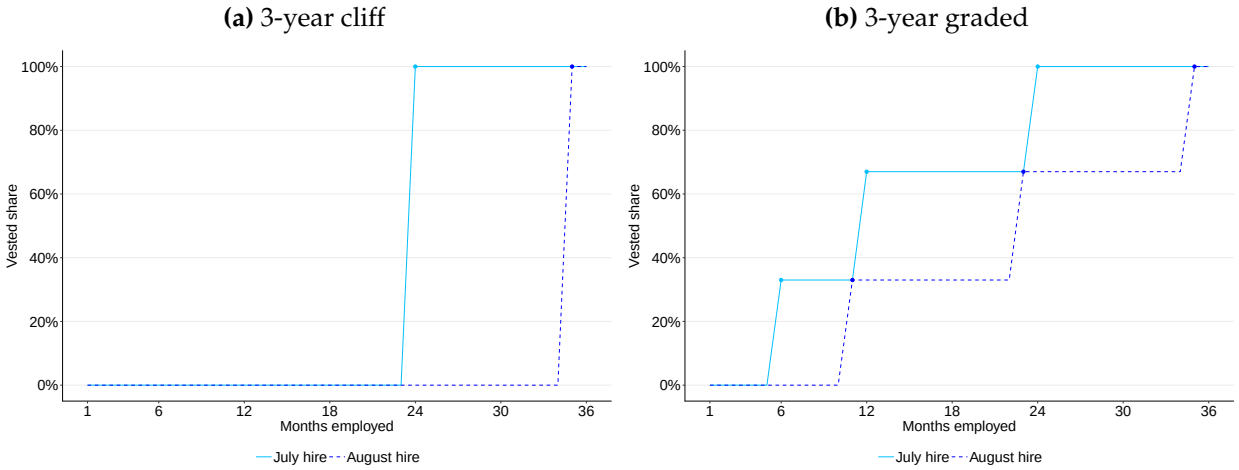
We identify 29 plans in our sample that credit service with the equivalency method. These plans use equivalence rules in which participants who are continuously employed for five months and at least one day of a sixth month in a given calendar year reach the 1,000-hour threshold and receive vesting credit.[26] For incumbent employees who were hired in a prior calendar year and are employed on January 1, vesting credit accrues when they reach their sixth month of calendar-year employment in June. For new employees hired in the current calendar year, vesting credit accrues six months after the hire date: January hires vest in June, February hires vest in July, March hires vest in August, and so on. The last group of new hires who receive credit for their first calendar year of employment are those who are hired in July and vest in December. Employees hired in August or later do not receive credit for their first calendar year of employment – their first vesting event occurs in June of their second calendar year. Appendix Figure A.10 illustrates the empirical operation of the equivalency method by plotting the share of vesting events (as reported in the administrative data) occurring in each calendar month. As expected, incumbents vest in June and new employees hired in July or earlier vest six months after hire.

The discontinuous treatment of July and August hires gives the quasi-random variation in vesting requirements that we exploit in this section. Figure 10 illustrates the differential vesting requirements that apply to July and August hires at equivalency-method plans. We use three-year cliff and three-year graded schedules as examples, but similar logic applies to cliff and graded schedules of different lengths. Because July hires vest in December of their first calendar year, they accrue their second year of credited service in June of their

[25]For the federal regulation formally defining the equivalency method and listing the permissible equivalence relations between calendar periods and credited service hours, see 29 CFR § 2530.200b-3 (e).

[26]These equivalence relations include rules that credit i) 45 hours for each week in which a participant works at least one hour, ii) 90 hours for each bi-weekly payroll period in which a participant works at least one hour, or iii) 190 hours for each month in which a participant works at least one hour. Under all three of these rules, participants who are continuously employed for five months and at least one day of a sixth month in a given calendar year reach the requisite 1,000 hours of credited service.

Figure 10: Effective vesting schedules for July and August hires at plans using the equivalency method



Note: July and August hires face different effective vesting schedules because July hires receive vested service credit for their first calendar year of employment but August hires do not.

second calendar year and their third year of credited service in June of their third calendar year. As a result, they earn full ownership of employer contributions after 24 months and receive partial ownership grants at 6 and 12 months under the graded schedule. August hires do not earn full ownership until June of their fourth calendar year, which is their 35th month of employment. Under the graded schedule, August hires earn partial ownership grants at 11 and 23 months and therefore have a weakly lower ownership share than July hires at all times (the inequality is strict except in months 1-5, 11, 23, and 35 onward).

The upshot of Figure 10 is that equivalency-method plans provide a natural experiment that identifies the retention effect. By studying July and August hires, we can compare participants who are hired at the same plan at essentially the same time, receive the same non-401(k) compensation in expectation, and face the same non-vesting 401(k) plan rules (including employer contribution rules), but are effectively subject to different 401(k) vesting schedules. If vesting has a retention effect, then August hires' longer effective vesting schedules should cause them to have longer average tenures than July hires.

It is important to emphasize that service crediting at equivalency-method plans is based on equivalency hours, not on *actual* hours worked. Service-crediting methods based on actual hours (which are more common at plans with hourly workers) could potentially be subject to manipulation in which employees take additional shifts to cross the 1,000-

hour threshold or employers rearrange shifts to keep employees below the 1,000-hour threshold. The only manipulation possible at equivalency-method plans concerns hire dates: employers could theoretically shift hire dates after August 1 in order to subject more employees to longer vesting schedules. Appendix Figure A.11 shows that such hire-date manipulation does not appear to occur in our sample, at least along observable dimensions. A roughly equal share of participants are hired in July and August; average age and income are also roughly flat across the cutoff.

We view the quasi-random vesting variation at equivalency-method plans as complementary to the cross-plan empirical strategy that we apply in Section 5. The relative strength of the equivalency-method strategy is its cleaner within-plan identification. The relative weaknesses are threefold. First, the equivalency-method sample is substantially smaller (see Table 1). Second, service-crediting rules may not be as salient as the vesting schedule itself; participants are more likely to know whether their plan has vesting requirements than they are to know whether their plan credits service using the elapsed-time or equivalency method. Third, comparing July and August hires at equivalency-method plans gives intensive-margin variation in the length of the vesting period, while the cross-plan approach in Section 5 gives more substantial extensive-margin variation in immediate versus non-immediate vesting. The consistency of our null retention results under both identification strategies gives us greater confidence in our conclusions.

6.2 Results

We begin by comparing raw survival probabilities for participants hired in different calendar months. If a retention effect is present, then August hires should be (discontinuously) more likely to remain employed beyond a given tenure length T. As in Section 5, we drop explicitly identified involuntary separations.

We take two different approaches to calculating survival probabilities. First, under the simpler approach, we restrict the data to participants who are hired sufficiently early in our sample period to observe the relevant survival outcome directly. For example,

since our data run through 2023,[27] when calculating observed survival probabilities for a tenure threshold of $T = 72$ months, we restrict to participants hired in 2017 or earlier. Using directly observable survival outcomes requires fewer assumptions but forces us to drop participants hired later in our sample period. Under the second approach, we estimate Kaplan-Meier survival probabilities from nonparametric separation hazard rates. In particular, let $m \in \{-6, -5, ..., 1\}$ be participants' hire month relative to the vesting cutoff, with $m = 0$ representing July hires and $m = 1$ representing August hires. For each relative hire month and each month of employment $t \in \{1, 2, ..., 72\}$, we estimate discrete-time separation hazard rates h_{mt} (i.e., the share of participants still employed in month t who separate in month t). For each survival threshold $T \in \{12, 24, 36, 48, 60, 72\}$, we infer the survival probability of participants hired in relative month m as $\prod_{t=1}^{T}(1 - h_{mt})$. This second approach makes better use of the available data (participants hired in later years can still be used to estimate hazard rates h_{mt} for small values of t) but requires more assumptions on the data-generating process.

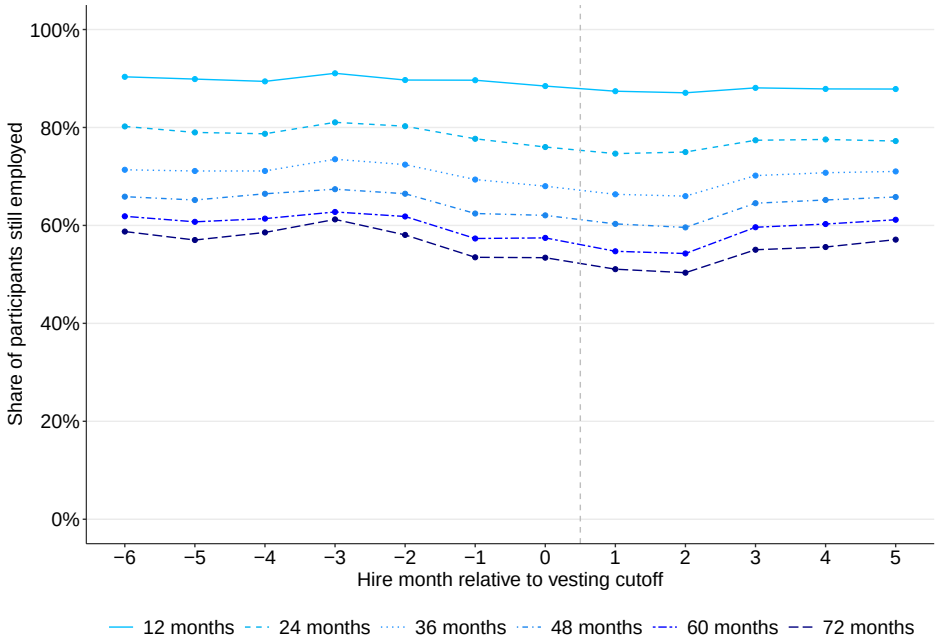
Figure 11 plots survival probabilities by relative hire month, with panel (a) showing directly observed probabilities and panel (b) showing probabilities inferred from separation hazard rates. Because average employment tenures trend downward during our sample period (see Table 2), the survival probabilities in panel (a) (which give more weight to participants hired early in the sample period) are shifted upward relative to panel (b). Despite the shift in levels, both plots tell the same story: there is no evidence of a retention effect when comparing July and August hires. Survival probabilities for each tenure threshold T are smooth across the July-August discontinuity. In fact, all 12 lines show a slight decrease between July and August, which is the opposite of a retention effect.

We next make more formal comparisons of survival probabilities that control for participant-level characteristics and plan fixed effects. In particular, for each survival threshold $T \in \{12, 24, 36, 48, 60, 72\}$, we restrict to participants hired in July or August

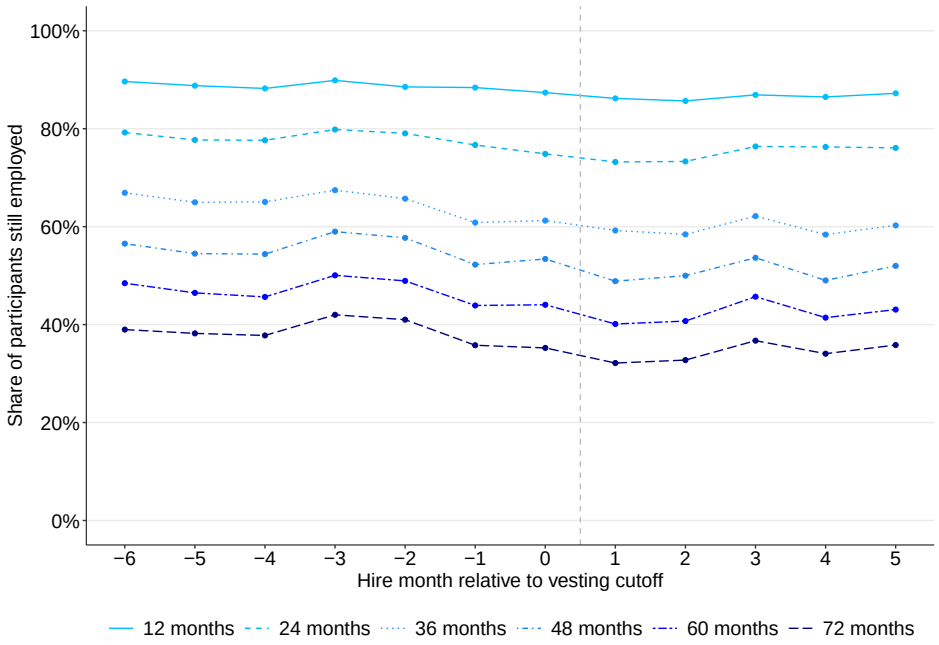
[27] As discussed in Section 3.1, we consider participants hired between 2010 and 2022 but have complete administrative data through 2023, which allows us to observe at least one year of post-hire data for all participants in our sample.

Figure 11: Survival probabilities by hire month

(a) Observed probabilities from uncensored participants

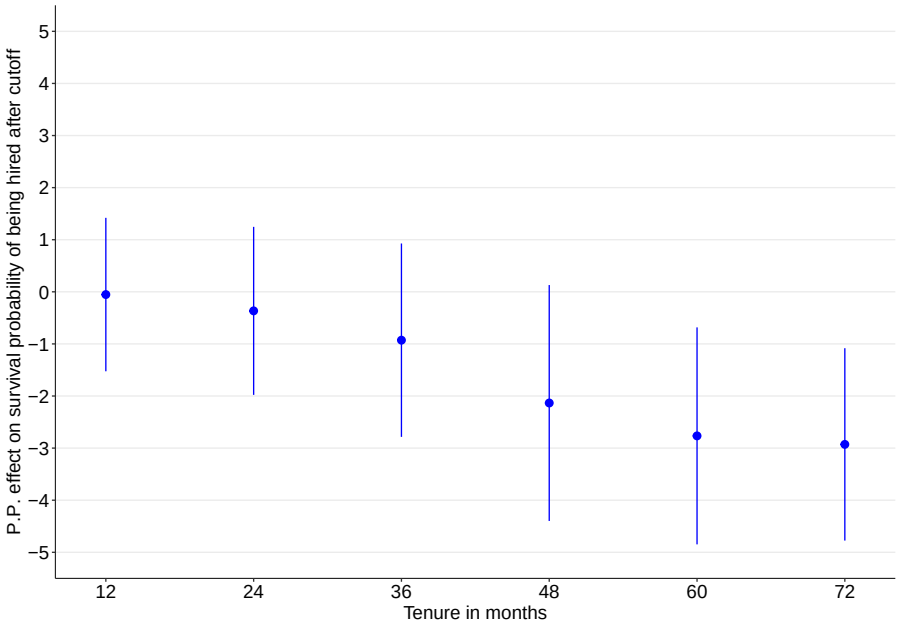


(b) Inferred probabilities from nonparametric hazard rates



Note: Participants hired in relative month 0 (July) or earlier receive vesting credit for their first calendar year of employment. Participants hired in relative month 1 (August) or later do not receive credit for their first calendar year and are thus subject to a longer effective vesting schedule. Panel (a) plots observed survival probabilities, considering only participants who are hired sufficiently early to observe the survival outcome (e.g., since our data run through 2023, we can only observe 72-month survival probabilities for participants hired in 2017 or earlier). Panel (b) considers all participants and shows Kaplan-Meier survival probabilities obtained from nonparametric separation hazard rates. Concretely, we estimate monthly separation hazard rates h_{mt} for each relative hire month $m \in \{-6, -5, ..., 5\}$ and month of employment $t \in \{1, 2, ..., 72\}$. For a given relative hire month m and survival threshold $T \in \{12, 24, 36, 48, 60, 72\}$, we infer the survival probability as $\prod_{t=1}^{T}(1 - h_{mt})$.

Figure 12: The survival effect of being hired after the vesting cutoff



Note: For each survival threshold $T \in \{12, 24, 36, 48, 60, 72\}$, we estimate a linear model where the observation is a participant, the dependent variable is an indicator for the participant remaining employed for at least T months, and the independent variables include: i) an indicator for being hired after the vesting cutoff (i.e., in August or later) and therefore being subject to a longer effective vesting schedule, ii) initial age at hire, iii) log of initial annualized income at hire, iii) fixed effects for calendar year of hire, and iv) plan fixed effects. We plot the coefficient estimates for the hired-after-vesting-cutoff indicator, along with their 95% confidence intervals. Standard errors are clustered by plan. We only consider participants hired sufficiently early to observe the survival outcome at threshold T (e.g., since our data run through 2023, we can only observe 72-month survival outcomes for participants hired in 2017 or earlier).

and estimate linear regressions of the form:

$$\mathbb{1}(y_{ij} \geq T) = \beta \mathbb{1}(m_{ij} = 1) + \gamma X_{ij} + \alpha_j + \theta_k + \epsilon_{ij}, \tag{6}$$

where y_{ij} is the maximum observed employment tenure for participant i in plan j, m_{ij} is the participant's hire month relative to the vesting cutoff (with $m_{ij} = 1$ for participants who are hired in August and therefore subject to a longer effective vesting schedule), X_{ij} is a vector of participant covariates including a quadratic of age at hire and the log of annualized income at hire, α_j are plan fixed effects, and θ_k are fixed effects for calendar year of hire. The coefficient of interest β is positive if a retention effect is present and participants hired in August or later have higher survival probabilities. We estimate (6) by restricting to participants who are hired sufficiently early to observe the survival outcome at T.

Figure 12 plots coefficient estimates for β along with their 95% confidence intervals. Consistent with the raw probabilities from Figure 11, all point estimates in Figure 12 are

negative, indicating that participants hired in August or later have slightly lower survival probabilities despite facing longer vesting schedules. Most estimates are statistically indistinguishable from zero, and with confidence intervals of 2-4 percentage points we can rule out positive retention effects larger than 1 percentage point.

7 Mechanisms

Using two different identification strategies, we find a robust null result: 401(k) vesting does not exert a systematic retention effect. In this section, we consider the underlying mechanisms that may explain this null result. To anchor the discussion, it is helpful to sketch a very simple theoretical framework describing an unvested 401(k) participant's decision problem as she considers leaving her current firm.

7.1 Theoretical framework

Consider a three-period model:

- In period 0, the worker is not yet fully vested and must decide whether to leave her current employer for a competing job offer. The worker has accumulated ex-ante retirement wealth of W and forfeits an amount F if she switches jobs. Her wage compensation at her current employer is y_{stay}. If she quits and switches jobs she instead earns y_{switch}, with $y_{switch} > y_{stay}$. She consumes her wage compensation, with utility given by $u(\cdot)$.

- In period 1, the worker becomes fully vested in her original employer's 401(k) contributions. Because switching jobs no longer involves a forfeiture cost and the higher outside wage offer y_{switch} makes her strictly better off,[28] the worker switches jobs with certainty if she did not already quit in period 0.[29] She again derives utility from consuming her wage compensation.

[28] Our underlying assumption is that the worker's outside wage option remains constant before and after the vesting deadline. If the worker receives a particularly attractive, non-recurring job offer in period 0, this assumption might not be warranted.

[29] We could also imagine the worker bargaining for the higher wage y_{switch} at her original employer, since she no longer faces a job-switching cost.

- In period 2, the worker retires and consumes her accumulated retirement wealth, which compounds with interest rate i.

The utility from switching or staying in period 0 is given by:

$$U_{switch} = u(y_{switch}) + \beta u(y_{switch}) + \beta^2 u((1+i)^2(W-F)), \tag{7}$$

$$U_{stay} = u(y_{stay}) + \beta u(y_{switch}) + \beta^2 u((1+i)^2 W). \tag{8}$$

The worker therefore switches in period 0 (and suffers a forfeiture) if:

$$U_{switch} - U_{stay} > 0 \tag{9}$$

$$u(y_{switch}) - u(y_{stay}) > \beta^2 [u((1+i)^2 W) - u((1+i)^2(W-F))]. \tag{10}$$

Using a linear approximation and assuming for simplicity that $\beta = 1/(1+i)$, we have:

$$u'(y_{stay}) \cdot (y_{switch} - y_{stay}) > u'((1+i)^2 W) \cdot F. \tag{11}$$

Since the worker can always choose to separate after becoming fully vested, she compares the one-period utility gain from a higher wage in period 0 with the loss of retirement utility from the vesting forfeiture.

Factors working against retention effect: Though highly stylized, this framework identifies two main factors that may explain our null retention result.

1. **Large wage gains at next job:** Participants may not respond to 401(k) vesting incentives if forfeitures are small relative to the wage gains realized from switching jobs. As the wage differential $y_{switch} - y_{stay}$ increases, workers become more willing to pay the forfeiture cost F and the retention power of vesting diminishes.[30]

2. **Informational or cognitive frictions:** Workers may be unaware that their 401(k) plans

[30] Our model only considers consumption utility. Workers may also separate for non-pecuniary reasons (e.g., family care responsibilities). While such events could contribute to our null retention finding, we lack sufficient data to analyze their impact.

include vesting provisions. Even if they are aware of vesting rules, they may struggle to understand and calculate the financial cost of quitting during the vesting period. Given low average levels of financial literacy among the U.S. workforce (Lusardi and Mitchell, 2023), knowledge gaps may play a significant role. Under-informed workers may perceive forfeiture costs as smaller than they actually are, even when the true forfeiture amount equals or exceeds the wage gain from switching jobs.

We directly test both explanations. First, using data on workers whose old and new jobs both have Vanguard-administered 401(k) plans, we compare the actual wage gains from job switching against the actual forfeiture losses. Second, we directly measure informational and cognitive frictions among 401(k) participants through a custom survey.

7.2 Wage gains among observable job switchers

For our main sample of job separators in Section 4, we do not generally observe the wage gain at the next job.[31] However, a subset of workers in our sample move to a new job at a different Vanguard-administered plan, making realized wage changes observable. Though it is considerably smaller than the main sample and only captures worker movement among a contained set of firms, the job-switcher sample allows us to make direct comparisons between workers' income gains and forfeiture losses.

To construct the job-switcher sample, we begin with the sample of separations from Section 4 and take the following filtering steps. First, to ensure that we observe wage gains between consecutive jobs, the worker must be hired at a different firm with a Vanguard-administered 401(k) plan within six months of their preceding separation date. Second, both the old and new plans must have income data (we continue to use directly reported W-2 income data if available and inferred income otherwise; see again Section 3.1). Finally, to reduce extrapolation noise, we use annualized income from the last (first) calendar year at the old (new) job with at least three months of employment. For example, if a worker

[31] Even without data on wage gains, (11) implies that retention effects should be larger for workers with larger potential forfeiture losses F. But in Figures 7c-7d, we show that retention effects are not stronger for workers with larger employer contribution balances (whether measured in levels or as a share of income). This is another indication that informational frictions have an explanatory role.

Table 4: Summary statistics for subsample of observable job switches

	Overall	By forfeiture outcome at old job	
		No forfeiture	Forfeiture
Number of job switches	66,729	50,870 (76.2%)	15,859 (23.8%)
Annual income at old job			
25th percentile	46,248	50,167	37,586
Median	76,267	82,792	59,729
75th percentile	127,658	138,436	94,429
Annual income at new job			
25th percentile	49,978	53,166	42,650
Median	83,235	88,463	69,340
75th percentile	135,086	142,452	110,312
Annualized income gain at new job			
Median	5.6%	3.5%	13.1%
Mean	23.3%	21.1%	30.5%
Forfeiture as share of old-job income			
Median			2.5%
Mean			4.3%

Note: See text for a description of the job-switcher sample. The annualized income gain is expressed as a share of old-job income (i.e., new minus old income divded by old income).

separates in March 2015 then we use annualized 2015 income as our measure of old-job wages; if she separates in February 2015 we instead use annualized 2014 income.[32]

Table 4 gives summary statistics for the set of roughly 67,000 observable job changes. Because higher-income plans are overrepresented in the job-switcher sample, annual incomes are slightly higher than in the main sample of separators. The median income gain is 5.6%, which is comparable to the 4-5% median wage gains for job switchers reported by the Atlanta Fed Wage Growth Tracker for most of the 2010-2022 period. A long right tail of large income gains (driven in part by large job-switching gains in the tight post-pandemic labor market) makes average gains significantly larger than median gains.[33] In raw terms, those who forfeited 401(k) balances at the old job have substantially higher income gains (median of 13.1% versus 3.5% for those without forfeitures). Forfeiture

[32]Similarly, we use the first calendar year at the new job if the hire date is in October or earlier and the second calendar year otherwise. If the worker never has a calendar year with at least three months of employment in either the old or new job, we drop her from the sample entirely. Because most of our income data comes from calendar-year W-2 figures, we cannot generally compute specific income figures for the 12-month period preceding separation or the 12-month period following hire.

[33]See Appendix Figure A.12, which plots wage gains among our job-switcher sample by year of separation.

Table 5: Regression analysis of wage gains for observable job switchers

	Dependent variable: proportional income gain at new job		
	(1)	(2)	(3)
1{Forfeiture at old job}	0.102	0.048	-0.034
	(0.008)	(0.009)	(0.008)
Controls			
Year-of-separation FE	X	X	X
Age and tenure at separation		X	X
Log annual income at old job			X

Note: See text for a description of the job-switcher sample. The dependent variable is the proportional income gain (new minus old income, divided by old income). Columns 2-3 control for quadratics of age and tenure at separation. Column 3 additionally controls for the log of annual income at the old job. The observation count for all specifications is 66,729, as in the first row of Table 4. Standard errors are in parentheses.

amounts are generally smaller than realized income gains: as a share of old-job income, the median forfeiture is 2.5% of income.

The relationship between forfeiture losses and income gains becomes more nuanced once we control for joint determinants of separations and wage gains. Table 5 shows estimates from regressions where the proportional income gain at the new job is the dependent variable. The independent variable of interest is a dummy equal to one if the worker incurred a forfeiture when they left their old job. Conditional on age, employment tenure, and old-job income, those with forfeitures see new-job income gains that are 3.4 percentage points *smaller* than those without forfeitures.

Tables 4-5 together tell a story of infra-marginal losses rather than marginal decision-making. Job-changers with forfeitures see larger average income gains, but only because these workers tend to be in early-career, lower-income jobs where large wage gains are already common. Once we control for observables, workers with forfeitures actually realize smaller wage gains. If optimizing workers were comparing wage gains to forfeitures in a framework like (11), then unvested workers would earn higher income gains to compensate for their forfeiture losses. The job-switcher results suggest that other factors, such as informational frictions, limit the marginal effect of forfeiture losses on workers' quitting behavior.

7.3 Informational frictions and participant survey

To test the informational frictions hypothesis, we use data from a Vanguard survey of its defined-contribution plan participants. The survey was fielded in October 2024 and yielded 1,018 complete responses from 100,000 email invitations, implying a response rate of about 1%. Each survey invitation was associated with a unique client identifier, allowing us to link the responses with plan- and participant-level information in the administrative data. See Appendix C for details on the survey sample and questionnaire. The invitee and respondent samples are roughly balanced in terms of income and plan-level vesting rules, but older and longer-tenured invitees were more likely to complete the survey.

Table 6 summarizes responses to a series of questions assessing participants' vesting knowledge and provides strong evidence of informational gaps. When asked to identify the definition of a 401(k) vesting schedule from a multiple-choice list, 31% of respondents supply the correct answer; a plurality of 47% answer "I do not know / I am not sure." The survey then provides respondents with the definition of a vesting schedule and asks in a yes-no question whether their current plan has vesting. Only 33% of respondents can say correctly whether their plan has a vesting schedule and the plurality (49%) again state that they do not know. Among respondents who correctly state that their plan has vesting, only 20% can accurately identify the length of their vesting schedule.

Disaggregating the results, we find substantial heterogeneity in participants' vesting knowledge. In particular, respondents with higher income and higher levels of demonstrated financial literacy are more likely to answer the knowledge-assessment questions correctly. Among respondents with income below $60,000, 15% provide the correct definition of vesting, 25% know whether their plan has a vesting schedule, and 27% correctly calculate forfeiture amounts in each of two vignettes concerning hypothetical separation scenarios. Among those with income above $120,000, 49% provide the correct definition, 41% know their plan's vesting policy, and 59% answer both vignettes correctly. When splitting respondents by the number of correct answers they supply to a standard set of five financial literacy questions, the gradient is even sharper: among those answering all five literacy questions correctly, 61% know the definition of vesting, 50% know whether

Table 6: Vesting knowledge in survey of current participants

		Share with correct response				
		All respondents			Correctly state their plan has vesting	
		(1)	(2)	(3)	(4)	(5)
	Number of respondents	Definition of vesting schedule	Forfeiture vignettes	Vesting at your plan - yes/no	Vesting at your plan - length	Are you fully vested
Overall	1,018	31%	45%	33%	20%	81%
Age						
20-29	21	10%	43%	38%	0%	0%
30-39	81	26%	46%	23%	38%	77%
40-49	178	28%	45%	25%	27%	73%
50-59	380	33%	50%	36%	23%	78%
≥ 60	358	34%	38%	36%	14%	89%
Income						
$< 30{,}000$	46	17%	24%	20%	17%	50%
30,000-59,999	208	14%	27%	26%	16%	64%
60,000-89,999	216	24%	43%	29%	21%	81%
90,000-119,999	167	40%	53%	39%	21%	88%
$\geq 120{,}000$	312	49%	59%	41%	25%	91%
Financial literacy						
0-2 correct	391	10%	22%	21%	16%	67%
3-4 correct	478	39%	55%	37%	23%	86%
All 5 correct	149	61%	71%	50%	20%	88%
Tenure						
0-2	239	27%	39%	28%	33%	45%
3-4	101	29%	39%	35%	19%	43%
5-6	83	31%	44%	29%	31%	100%
7-10	126	33%	43%	31%	17%	90%
≥ 11	469	34%	49%	36%	15%	95%
Vesting status (vesting plans)						
Not fully vested	136	31%	35%	37%	30%	40%
Fully vested	435	31%	46%	40%	17%	93%

Note: For all questions assessing vesting knowledge, the following shares add to 100%: i) the share answering correctly (shown above), ii) the share answering incorrectly, and iii) the share answering "I do not know/I am not sure." We obtain age and tenure information and plan-level vesting rules from the administrative data. If a participant's plan has more than one vesting schedule (which can occur if the employer makes both matching and nonmatching contributions), we associate the participant with the longest vesting schedule at their plan. The survey also asks respondents to base their answers on the longest vesting schedule at their plan. We use annualized 2023 income from the administrative data if available and respondents' self-reports otherwise. The survey asks respondents five standard financial literacy questions; the fourth panel splits respondents according to the number of literacy questions they answered correctly. Question (1) asks respondents to identify the definition of a vesting schedule from a multiple-choice list. Question (2) comprises two vignettes in which respondents are asked to calculate forfeiture amounts in simple hypothetical separation scenarios (the correct forfeiture amount is zero in one vignette and nonzero in the other); we report the share of respondents answering both vignettes correctly. The vignettes were shown after respondents were provided with the definition and an example of a vesting schedule. Question (3) asks respondents whether their current plan has a vesting schedule and was posed after respondents were provided with the definition and an example of a vesting schedule. Question (4) asks respondents for the length of their current plan's vesting schedule and was only shown to respondents who indicate that their plan has a vesting schedule. Question (5) asks respondents if they are currently fully vested in their employer contributions and was only shown to respondents who indicate that their plan has a vesting schedule. For purposes of this table, we only consider responses to Questions (4) and (5) if the respondent *correctly* indicates that their plan has a vesting schedule. See Appendix C for a description of the survey sample and the survey questionnaire.

their plan has vesting, and 71% answer the vignettes correctly.

Perhaps more positively, among respondents who know that their plan has vesting, 81% can state correctly whether they are currently fully vested in their employer contributions.[34] The share answering correctly remains above 50% even among respondents with lower incomes and lower levels of financial literacy. Splitting the sample by employment tenure provides an explanation for these outlying results: the high share of correct responses is driven entirely by respondents with over five years of tenure. We hypothesize that long-tenured participants know they have enough tenure to be fully vested under any reasonable vesting schedule and can answer this question correctly without knowing the precise length of their schedule. Splitting by binary vesting status instead of tenure shows more clearly that knowledge is lowest among participants for whom it is potentially most relevant: 40% of non-fully vested respondents know whether they are fully vested, compared to 93% of fully vested respondents.

In conclusion, the survey reveals significant informational frictions among participants. Only about a third of survey respondents know what a vesting schedule is and whether their current plan has one. To put these results into context, the survey also asks respondents about non-vesting rules at their plan. Responses to the non-vesting questions are more accurate: 57% of respondents correctly identify the type of employer contributions they receive (matching, nonmatching, both, or neither), 57% correctly state whether their plan allows loans, and 50% correctly state whether their plan allows hardship withdrawals. Participants appear to be less well informed about vesting schedules than they are about other 401(k) features, perhaps because contribution and withdrawal rules are relevant during employment while vesting rules are enforced only at separation.

8 Conclusion

We provide evidence on the empirical effects of 401(k) vesting requirements. Vesting has received less research attention than other aspects of the defined-contribution system but

[34]This question was only shown to respondents who indicate that their plan has a vesting schedule. As with the question about the length of participants' vesting schedules, we further restrict to respondents who *correctly* indicate that their plan has a vesting schedule.

is a relevant consideration for many plan participants, with forfeitures occurring in 30% of separations. We show that these forfeitures occur disproportionately among lower-income workers and find no evidence that vesting exerts a causal retention effect. Informational frictions appear to be an important explanation for the lack of retention effect, as most respondents in a survey of 401(k) participants are unaware of vesting requirements and do not know their plan's vesting rules.

Given these results, a promising opportunity for future research is to study the optimal design of 401(k) vesting regulations, taking into account our null retention finding and considering possible general-equilibrium effects on other outcomes like wages and employment. More broadly, future work should also explore how informational frictions affect the use of other non-wage benefits and potentially contribute to inequality in total compensation.

References

ADP RESEARCH (2023): "The Hidden Truth About Promotions?" https://www.adpri.org/wp-content/uploads/2023/10/TaW_Q32023v2.pdf.

ALDATMAZ, S., P. OUIMET, AND E. D. V. WESEP (2018): "The Option to Quit: The Effect of Employee Stock Options on Turnover," *Journal of Financial Economics*, 127, 136–151.

BAKER, A. C., D. F. LARCKER, AND C. C. WANG (2022): "How Much Should We Trust Staggered Difference-in-Differences Estimates?" *Journal of Financial Economics*, 144, 370–395.

BALSAM, S., R. GIFFORD, AND S. KIM (2007): "The Effect of Stock Option Grants on Voluntary Employee Turnover," *Review of Accounting and Finance*, 6, 5–14.

BIASI, B. (2019): "Higher Salaries or Higher Pensions? Inferring Preferences from Teachers' Retirement Behavior," SSRN Working Paper.

——— (2024): "Salaries, Pensions, and the Retention of Public-Sector Employees: Evidence from Wisconsin Teachers," Working Paper.

CENGIZ, D., A. DUBE, A. LINDNER, AND B. ZIPPERER (2019): "The Effect of Minimum Wages on Low-Wage jobs," *The Quarterly Journal of Economics*, 134, 1405–1454.

CHOI, J. J., D. LAIBSON, J. CAMMAROTA, R. LOMBARDO, AND J. BESHEARS (2024): "Smaller than We Thought? The Effect of Automatic Savings Policies," Working Paper.

CHOUKHMANE, T., J. COLMENARES, C. O'DEA, J. ROTHBAUM, AND L. D. SCHMIDT (2023): "Who Benefits from Retirement Saving Incentives in the U.S.? Evidence on Racial Wealth Gaps in Retirement Wealth Accumulation," Working Paper.

CLARK, J. W. (2023): "How America Saves 2023," Vanguard Report.

COLE, A. AND B. TASKA (2023): "Worker Valuation of Retirement Benefits," Working Paper.

COPELAND, C. (2023): "Trends in Employee Tenure, 1983-2022," https://www.ebri.org/content/trends-in-employee-tenure-1983-2022.

FEDERAL RESERVE BANK OF ATLANTA (2025): "Wage Growth Tracker," https://www.atlantafed.org/chcs/wage-growth-tracker; accessed on 09/02/25.

GONG, Q., H. ZHANG, AND L.-A. ZHOU (2023): "Retention Effects of Employee Stock Options: Evidence from Bunching at Vesting Dates," Working Paper.

GREIG, F., A. MADAMBA, G. CARRANZA, C. O'DEA, T. CHOUKHMANE, AND L. SCHMIDT (2024): "Are Employers Optimizing Their 401(k) Match?" https://papers.ssrn.com/sol3/papers.cfm?abstract_id=484777.

HOFFMANN, F. AND V. VLADIMIROV (2025): "Worker Runs," *Journal of Finance*, 80, 937–979.

HOU, W. AND G. SANZENBACHER (2020): "Measuring Racial/Ethnic Retirement Wealth Inequality," SSRN Working Paper.

JESZECK, C. A., T. CROSS, A. JACOBS, S. CHAPMAN, K. MORRIS, R. PATTERSON, AND S. SPENCE (2016): "401(k) Plans: Effects of Eligibility and Vesting Policies on Workers' Retirement Savings," SSRN Working Paper.

JOCHEM, T., T. LADIKA, AND Z. SAUTNER (2018): "The Retention Effects of Unvested Equity: Evidence from Accelerated Option Vesting," *Review of Financial Studies*, 31, 4142–4186.

JORDÀ, Ò., K. KNOLL, D. KUVSHINOV, M. SCHULARICK, AND A. M. TAYLOR (2019): "The Rate of Return on Everything, 1870–2015," *The Quarterly Journal of Economics*, 134, 1225–1298.

KIM, K. AND C. WROBLEWSKI (2025): "From Pensions to Personnel: The Incentive Effects of Retirement Benefits on Retention," Working Paper.

KOTLIKOFF, L. J. AND D. A. WISE (1987): "The Incentive Effects of Private Pension Plans," in *Issues in Pension Economics*, ed. by Z. Bodie, J. B. Shoven, and D. A. Wise, University of Chicago Press, 283–340.

LADIKA, T. AND S. SAUTNER (2020): "Managerial Short-Termism and Investment: Evidence from Accelerated Option Vesting," *Review of Finance*, 24, 305–344.

LUSARDI, A. AND O. S. MITCHELL (2023): "The Importance of Financial Literacy: Opening a New Field," *Journal of Economic Perspectives*, 37, 137–154.

NI, S. AND M. PODGURSKY (2016): "How Teachers Respond to Pension System Incentives: New Estimates and Policy Applications," *Journal of Labor Economics*, 34, 1075–1104.

OYER, P. AND S. SCHAEFER (2005): "Why Do Some Firms Give Stock Options to All Employees? An Empirical Examination of Alternative Theories," *Journal of Financial Economics*, 76, 99–133.

PRINCE, S. J., T. G. AZIZKHAN, C. R. PRINCE, AND L. GORMAN (2024): "The Effects of 401(k) Vesting Schedules - in Numbers," *Yale Law Journal*, 134.

SABELHAUS, J. AND A. H. VOLZ (2022): "Wealth Inequality and Retirement Preparedness: A Cross-Cohort Perspective," SSRN Working Paper.

SCHILLER, B. R. AND R. D. WEISS (1979): "The Impact of Private Pensions on Firm Attachment," *The Review of Economics and Statistics*, 61, 369–380.

SUN, Q. AND M. Z. XIAOLAN (2019): "Financing Intangible Capital," *Journal of Financial Economics*, 133, 564–588.

TAN, F., F. GREIG, A. S. CLARKE, K. KHANG, K. MCKINNON, AND V. ZHANG (2023): "The Vanguard Retirement Outlook: A National Perspective on Retirement Readiness," https://institutional.vanguard.com/content/dam/inst/iig-transformation/insights/pdf/2023/the-vanguard-retirement-outlook.pdf.

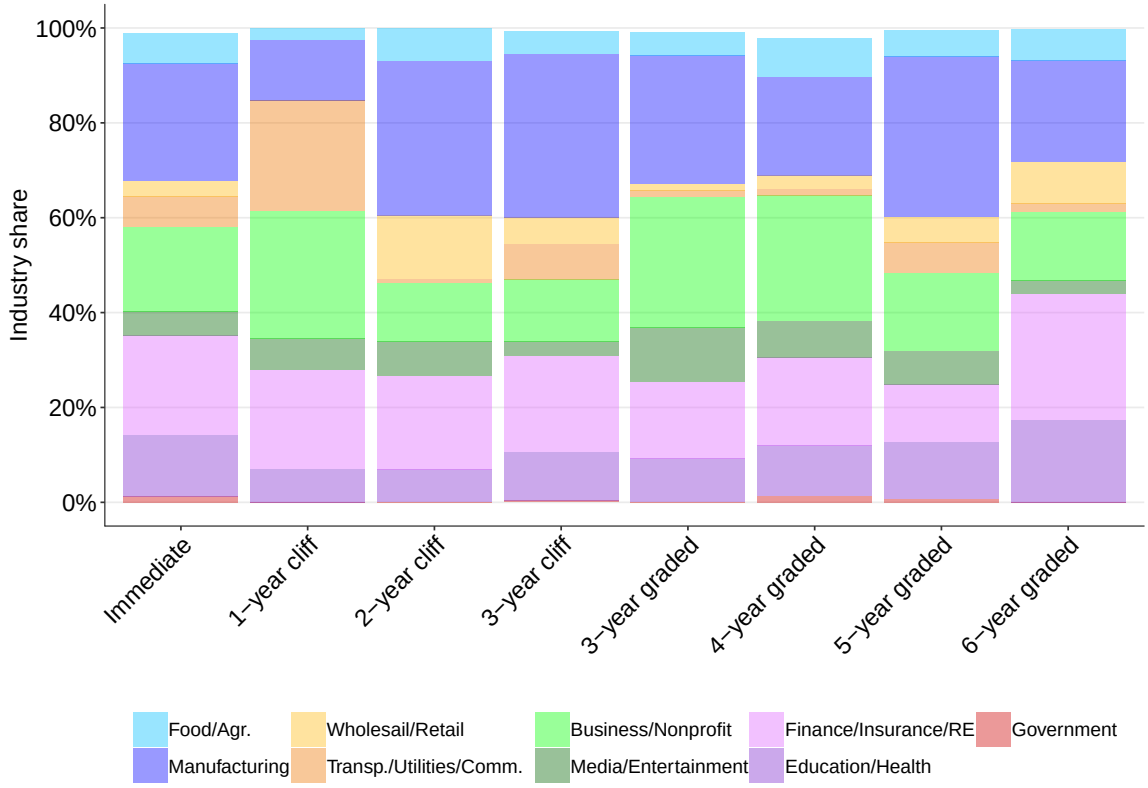
WARD, J. (2022): "The New Vesting Schedule Debate," *Planadviser*, https://www.planadviser.com/exclusives/new-vesting-schedule-debate/.

WOLFF, E. N. (2023): "Understanding Trends in Hispanic and African American Retirement Preparedness in the US," SSRN Working Paper.

Internet Appendix

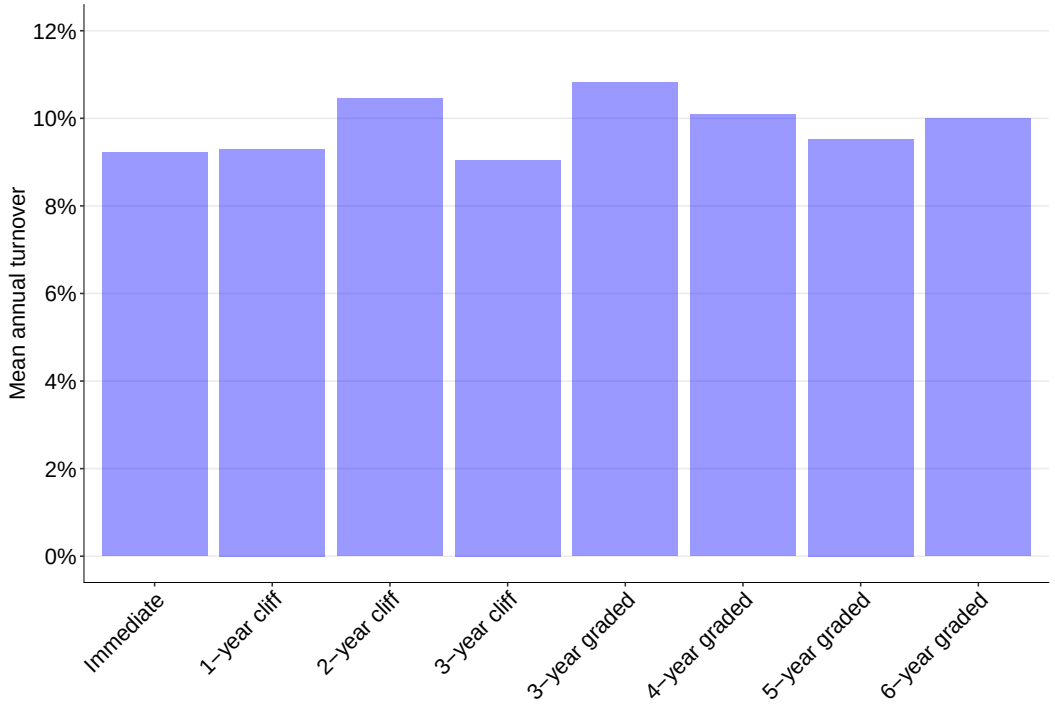
Appendix A Additional empirical results

Figure A.1: Vesting schedule prevalence by industry



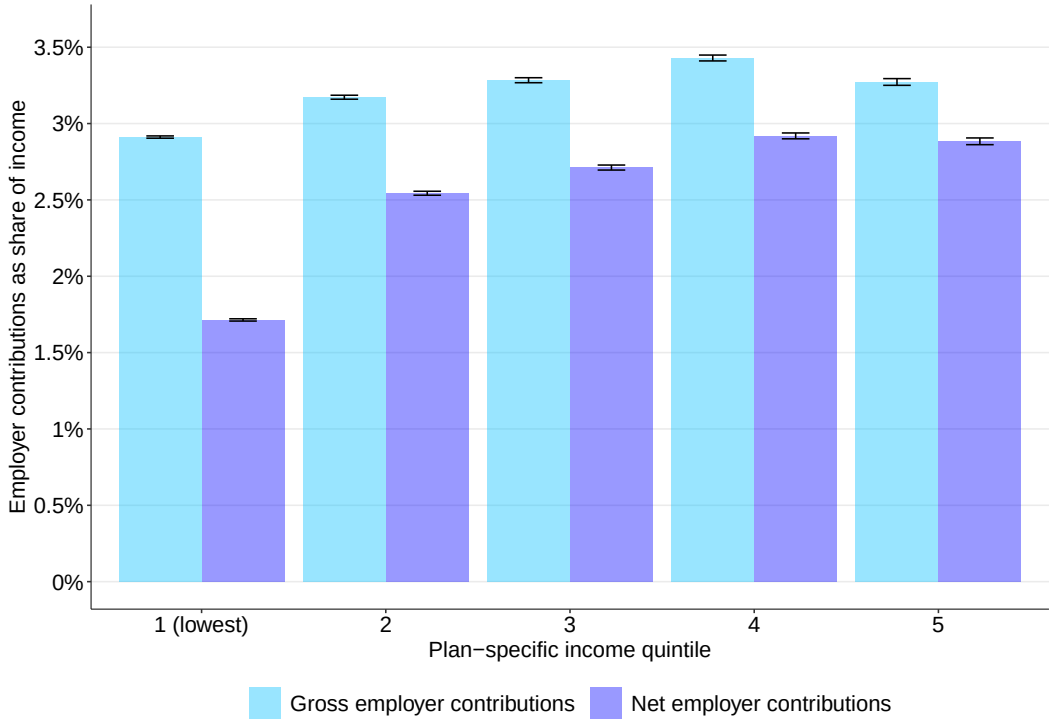
Note: Vesting schedules are as of 2022 and industry shares are plan-weighted. For plans with multiple vesting schedules, we take the schedule with the longest vesting period. The industry labels come directly from the recordkeeping data and are not assigned by the authors.

Figure A.2: Participant turnover by vesting schedule



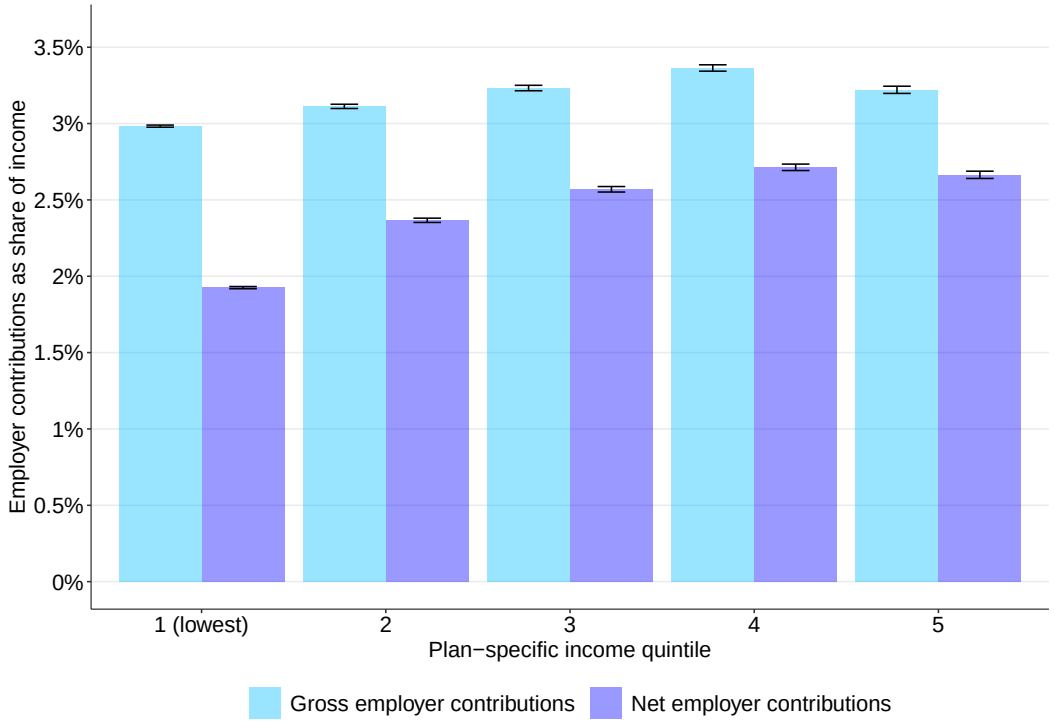
Note: For each plan-year, we calculate the turnover rate as the ratio of separations to the total number of participants employed at the plan in the same year. Each plan is assigned to a single vesting schedule; for plans with multiple vesting schedules, we take the schedule with the longest vesting period. The bar represents aggregates across our entire sample period.

Figure A.3: Distributional effects, classifying participants based on average annual income



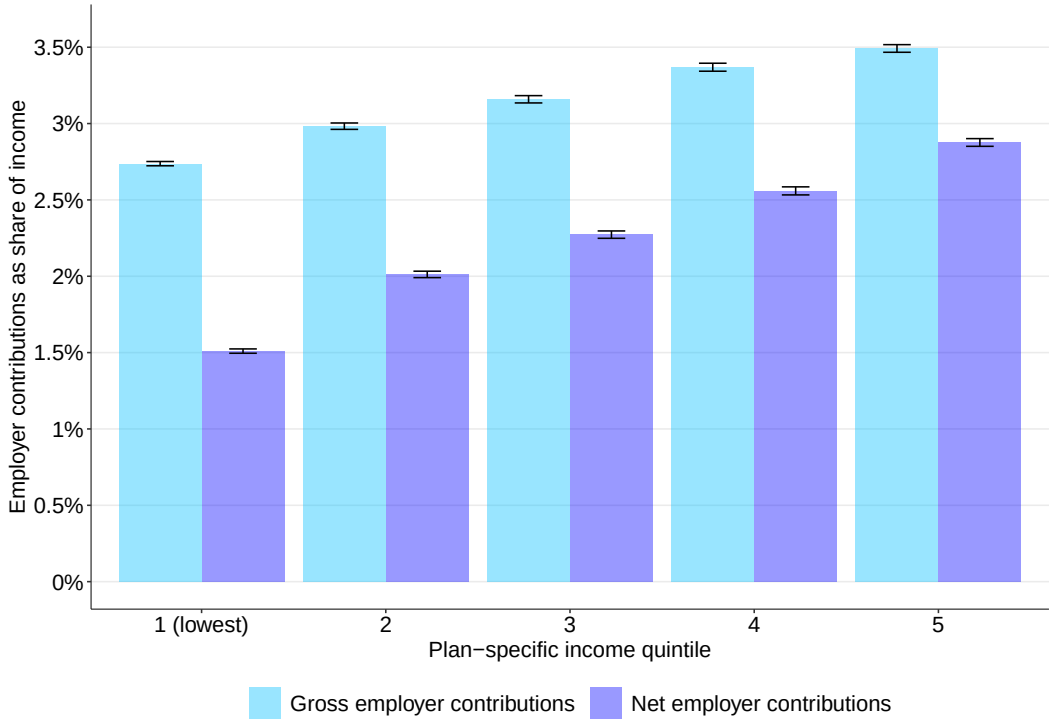
Note: This plot reproduces Figure 3, but sorts separating participants into income quintiles based on their average annual income at the plan rather than their final annual income.

Figure A.4: Distributional effects, classifying participants based on first annual income



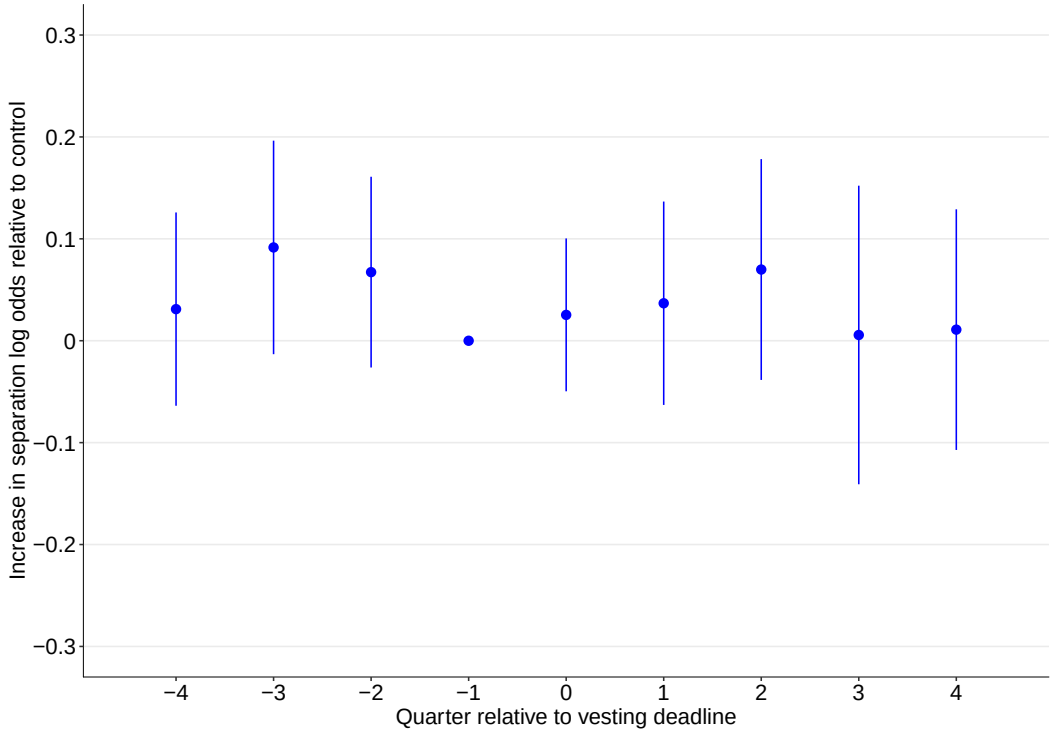
Note: This plot reproduces Figure 3, but sorts separating participants into income quintiles based on their first observable annual income at the plan rather than their final annual income.

Figure A.5: Distributional effects, restricting to separations in 2018 or earlier



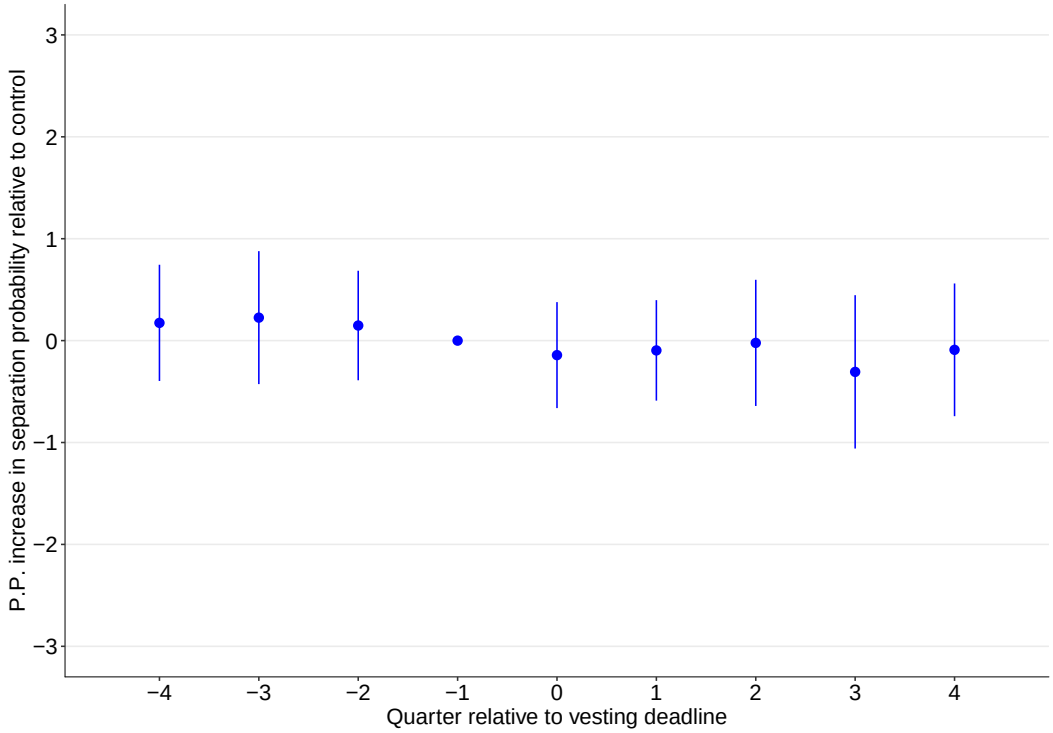
Note: This plot reproduces Figure 3, but restricts to participants separating in 2018 or earlier (the subset whose forfeiture transactions, which may occur with up to a 5-year lag after separation, are not time-censored in our data).

Figure A.6: Nonlinear event study for 3-year cliff vesting schedules



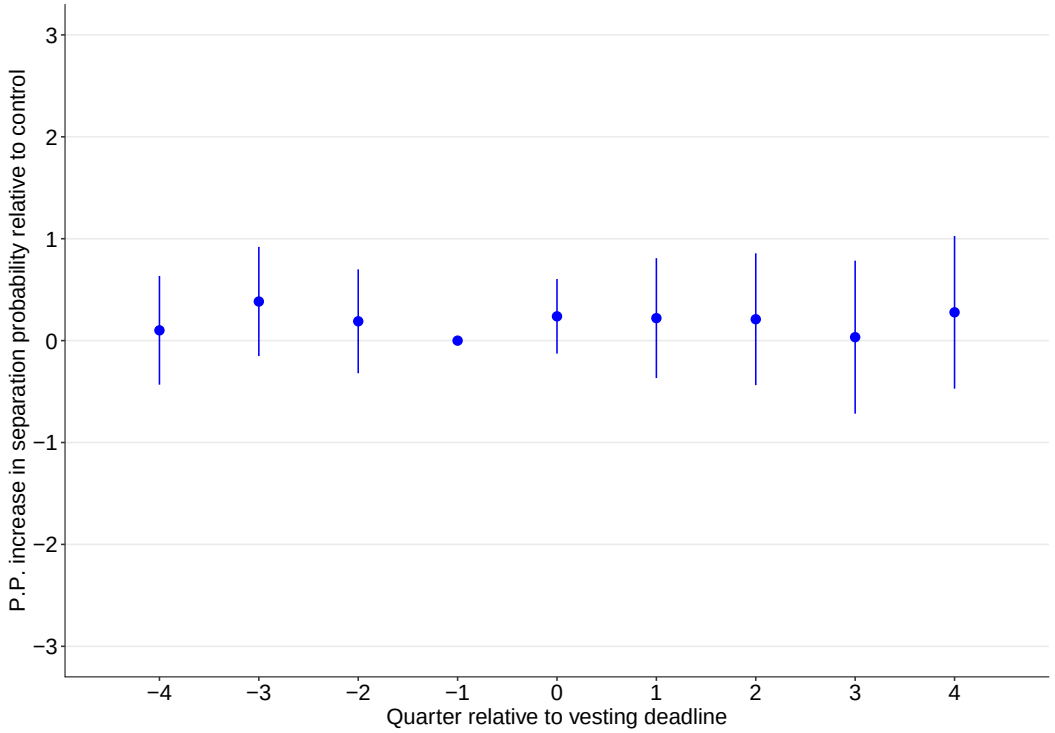
Note: We estimate a nonlinear difference-in-difference event study with participants in three-year cliff vesting plans as the treatment group and participants in immediate-vesting plans as the control group. We estimate a discrete survival hazard model using complementary log-log as the link function. We plot the coefficient estimates β_k from (4) along with their 95% confidence intervals. The coefficient estimates are normalized with respect to relative quarter $k = -1$. The outcome variable is an indicator equal to one if the participant voluntarily separates in a given quarter. We control for age and income-bin fixed effects, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are bootstrapped, not clustered at the plan-level.

Figure A.7: Event study for 3-year cliff vesting schedules, restricting to plans with a single type of employer contribution



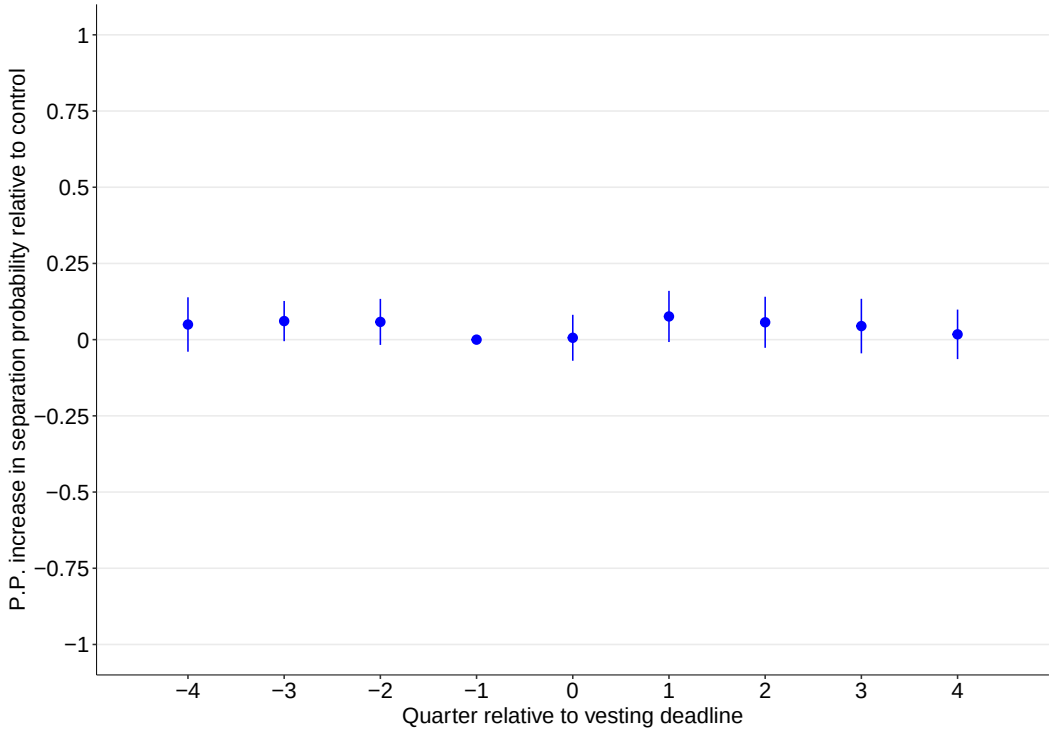
Note: We estimate a difference-in-difference event study with participants in three-year cliff vesting plans as the treatment group and participants in immediate-vesting plans as the control group. We subset to participants who were auto-enrolled into participation. We plot the coefficient estimates β_k from (4) along with their 95% confidence intervals. The coefficient estimates are normalized with respect to relative quarter $k = -1$. The outcome variable is an indicator equal to one if the participant voluntarily separates in a given quarter. We control for age and income-bin fixed effects, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are clustered at the plan level.

Figure A.8: Event study for 3-year cliff vesting schedules, restricting to automatic-enrollment plans



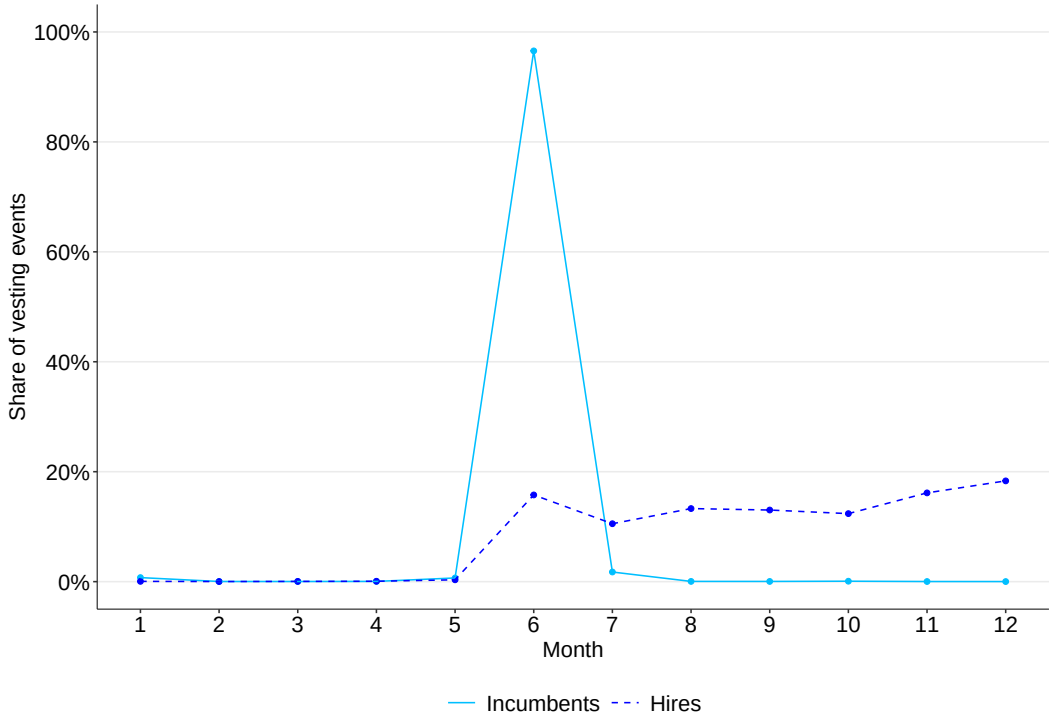
Note: We estimate a difference-in-difference event study with participants in three-year cliff vesting plans as the treatment group and participants in immediate-vesting plans as the control group. We subset to participants who were auto-enrolled into participation. We plot the coefficient estimates β_k from (4) along with their 95% confidence intervals. The coefficient estimates are normalized with respect to relative quarter $k = -1$. The outcome variable is an indicator equal to one if the participant voluntarily separates in a given quarter. We control for age and income-bin fixed effects, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are clustered at the plan level.

Figure A.9: Event study for 3-year cliff vesting schedules, restricting to involuntary separations



Note: We estimate a difference-in-difference event study with participants in three-year cliff vesting plans as the treatment group and participants in immediate-vesting plans as the control group. We restrict to separations that are explicitly coded in the recordkeeping data as involuntary (we drop these involuntary separations from the main analysis; see the discussion in Section 5.2). We plot the coefficient estimates β_k from (4) along with their 95% confidence intervals. The coefficient estimates are normalized with respect to relative quarter $k = -1$. The outcome variable is an indicator equal to one if the participant voluntarily separates in a given quarter. We control for age and income-bin fixed effects, calendar-year and quarter fixed effects, plan industry fixed effects, and plan-size-decile fixed effects. Standard errors are clustered at the plan level.

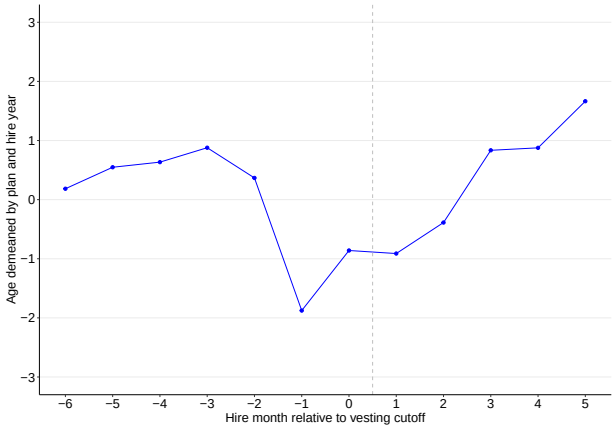
Figure A.10: Vesting events by calendar month at plans using the equivalency method



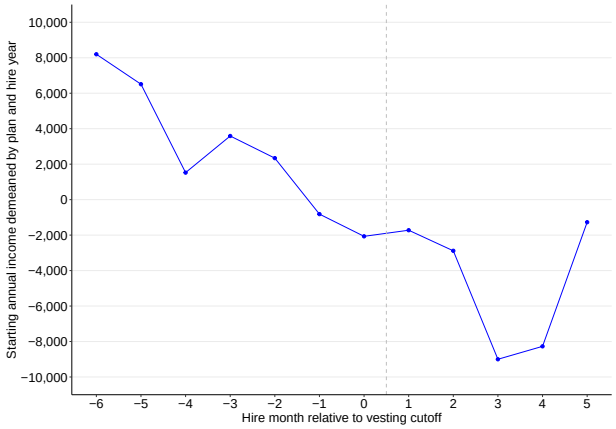
Note: A vesting event is the date (occurring once per calendar year) on which a participant's credited years of service increment by one. Incumbents are participants who were hired in a prior calendar year and hires are participants who are hired in the current calendar year.

Figure A.11: Covariates by hire month at plans using the equivalency method

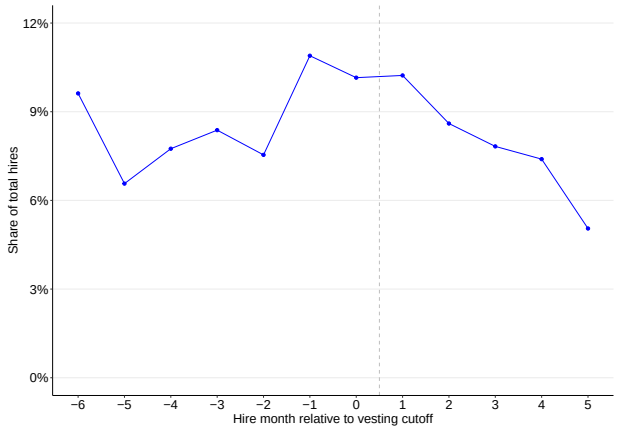
(a) Plan-demeaned age



(b) Plan-demeaned income

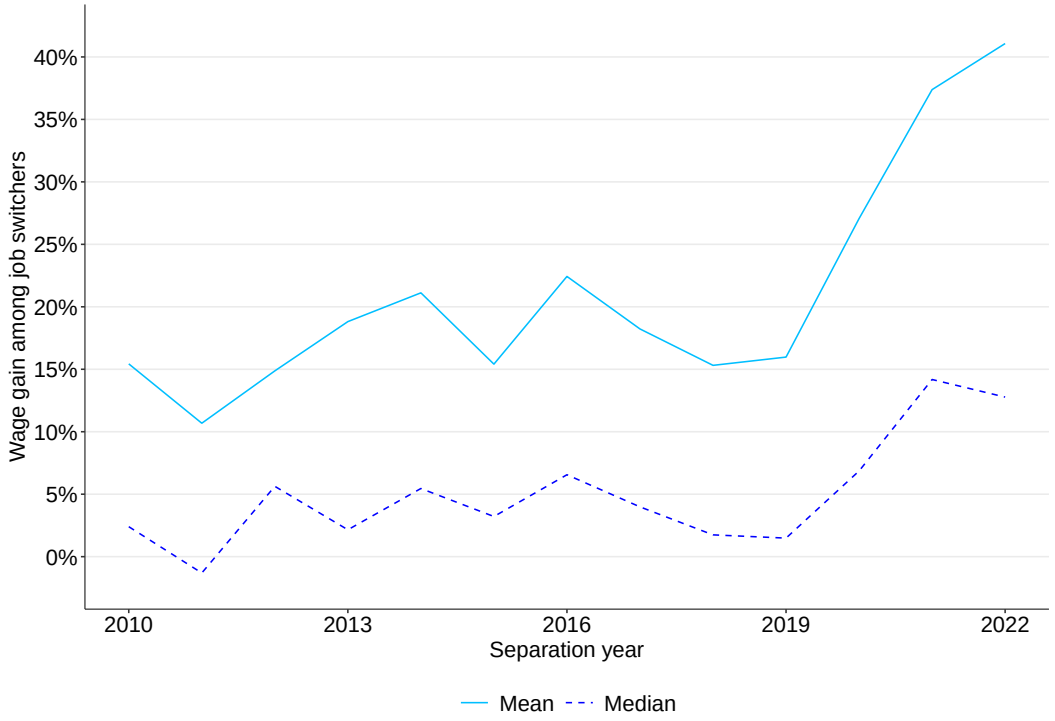


(c) Share of hires



Note: Participants hired in relative month 0 (July) or earlier receive vesting credit for their first calendar year of employment. Participants hired in relative month 1 (August) or later do not receive credit for their first calendar year and are thus subject to a longer effective vesting schedule. Panel (a) shows initial age at hire, demeaned by plan and hire year. Panel (b) shows annualized income in the first calendar year of employment, demeaned by plan and hire year. Panel (c) shows the raw distribution of hire dates.

Figure A.12: Wage gains among observable job switchers, by year



Note: We begin with the main sample of separations (see Section 4) and restrict to those who begin a new job at a different Vanguard-administered plan within six months of their separation date. We use annualized income from the last (first) calendar year at the old (new) job with at least three months of employment. For example, if a worker separates in March 2015 then we use annualized 2015 income as our measure of old-job wages; if she separates in Februrary 2015 we instead use annualized 2014 income. Similarly, we use the first calendar year at the new job if the hire date is in October or earlier and the second calendar year otherwise. The figure plots annualized income gains (new job minus old job) by year of separation.

Table A.1: Regression estimates of age and income disparities in forfeiture rates

Dependent variable: nonzero forfeiture	(1)	(2)
Age group (60+ omitted)		
18-30	0.257	0.212
	(0.017)	(0.012)
30-40	0.173	0.144
	(0.011)	(0.011)
40-50	0.139	0.115
	(0.009)	(0.009)
50-60	0.092	0.077
	(0.007)	(0.007)
Income quintile (top omitted)		
Bottom	0.181	0.147
	(0.039)	(0.041)
Second	0.151	0.116
	(0.034)	(0.0215)
Third	0.092	0.076
	(0.016)	(0.014)
Fourth	0.054	0.039
	(0.011)	(0.012)
Year fixed effects	X	X
Plan fixed effects		X

Note: We estimate linear regressions using the entire sample of separation-related withdrawals during the 2010-2022 period. The dependent variable is an indicator equal to one if the separation involves a nonzero forfeiture of employer contributions. The omitted age group is 60+ and the omitted income group is the top quintile. Standard errors are in partentheses and are clustered by plan.

Table A.2: Summary statistics for main D-i-D sample

	Treatment (3-year cliff)	Control (Immediate)
Number of plans	181	428
Number of participants	162,788	439,309
Median age	34	34
Income		
25th percentile	39,927	33,609
50th percentile	60,571	55,995
75th percentile	94,722	94,420
Participant-weighted industry shares		
Manufacturing	35%	34%
Business/Professional	27%	17%
Finance	17%	8%
Health/Education	7%	14%
Other	15%	27%

Note: Age and income are measured at the time of hire. Income figures are given in real 2022 terms. Industry categories are fixed at the plan level and are those assigned by the data provider.

Table A.3: D-i-D regression estimates for 3-year cliff schedule

Sample	Main	Unique Schedule	Auto-enrollment
Post	-0.2286	-0.2071	-0.3272
	(0.1064)	(0.1209)	(0.1767)
Treatment	-0.6338	-0.8761	-0.6680
	(0.2901)	(0.4428)	(0.3875)
Post x Treatment	-0.0763	-0.2745	0.0212
	(0.1289)	(0.1958)	(0.1993)
Controls	X	X	X
Observations	3,156,687	2,100,693	1,958,120

Note: This table shows the D-i-D regression estimation for the main specification described in Section 5 and similar robustness exercises. The specification used is described in Equation 4. The specification includes age quadratic, log income, calendar year, quarter, plan size-decile and industry FEs as controls. All coefficient estimates are multiplied by 100 to give effects in percentage-point terms. Standard errors are clustered by plan.

Table A.4: Heterogeneity of D-i-D regression estimates by schedule

	Cliff		Graded			
	2Y	3Y	3Y	4Y	5Y	6Y
Post x Treatment	-0.6871	-0.0763	-0.1197	-0.0372	-0.0255	0.2316
	(0.4406)	(0.1289)	(0.1272)	(0.1629)	(0.1009)	(0.1437)
Controls	X	X	X	X	X	X
Observations	3,504,151	3,156,687	11,042,164	6,442,298	3,367,946	1,263,459

Note: This table shows the D-i-D regression estimation for the specification described in Section 5 for cliff and graded schedules. The specification used is for cliff schedules is Equation 4 and the one used for graded schedules is Equation 5. The specifications include age quadratic, log income, calendar year, quarter, plan size-decile and industry FEs as controls. All coefficient estimates are multiplied by 100 to give effects in percentage point terms. Standard errors are clustered by plan.

Table A.5: Survival regression estimates for equivalency-method plans

Survival horizon (months)	12	24	36	48	60	72
Hired after vesting cutoff	-0.053	-0.366	-0.928	-2.134	-2.765	-2.929
	(0.751)	(0.823)	(0.947)	(1.155)	(1.062)	(0.942)
Plan fixed effects	X	X	X	X	X	X
Year-of-hire fixed effects	X	X	X	X	X	X
Age and income controls	X	X	X	X	X	X
Observations	15,369	15335	13,852	12,964	11,490	9,862

Note: This table shows the underlying survival regression estimates for equivalency-method plans that reflect the specification in (6) and are summarized in Figure 12. The specification controls for a quadratic of age at hire and the log of annualized income at hire. All coefficient estimates are multiplied by 100 to give effects in percentage-point terms. Standard errors are clustered by plan.

Appendix B Lifetime simulation calibration

In this appendix, we describe the simulation process and the inputs for the lifetime simulation model in Section 4.3. We estimate all inputs (besides historical asset returns) from the administrative recordkeeping data, ensuring that our model calibration is an internally consistent representation of 401(k) plan participants. The simulation proceeds as follows:

- The worker earns age-specific income from a deterministic lifetime income path and participates in her 401(k) plan according to an age- and income-specific participation rate. If she participates, she saves according to an age- and income-specific contribution rate and receives matching contributions according to her plan's match formula. She also receives nonmatching contributions (which are independent of her participation and contribution decisions) if her employer provides them. Her 401(k) wealth grows according to market returns.
- The worker separates stochastically according to an age-, income-, and tenure-specific separation hazard rate. If separation occurs and she is not fully vested, she suffers a forfeiture.
- The worker's age increases by one year and the process repeats.
- The worker's career ends at age 59 and her 401(k) wealth earns returns until age 65.

The worker's annual income proceeds along a deterministic lifetime path. We construct quintile-specific paths using the cross-sectional income distribution among our balanced panel of plans in 2022. Specifically, for each age between 25 and 59, we split employees of that age into quintiles and compute average annual income within each quintile.[35] Since we now model transitions between plans, we define the quintiles on a cross-plan (rather than within-plan) basis. Figure B.1 shows the resulting quintile-specific paths. We assume that workers belong to a single income quintile for the duration of their careers and thus do not switch between the distinct income paths shown in Figure B.1.

[35]To limit the influence of right-tail outliers, we use median rather than average income for the top quintile.

To simulate market returns, we assume that all of the worker's 401(k) contributions are invested in a target date fund and that real equity and bond returns are equal to the historical U.S. averages in the database compiled by Jordà et al. (2019).[36] Because we construct income paths from the 2022 income distribution and use real asset returns, all simulated dollar values are in real 2022 terms.

When the worker begins a new job (at age 25 and after subsequent separations), we sample a plan from our balanced panel. For the duration of the corresponding job spell, the worker earns the matching and nonmatching contributions – and is subject to the associated matching and nonmatching vesting schedules – that were in force at the sampled plan in 2022. Our goal is to simulate a worker moving through the current distribution of employer contributions and vesting schedules as her career progresses. We use plan sampling weights that are specific to each income quintile (reflecting the income composition of each plan in 2022), so our job-transition process captures observed sorting patterns between worker income and plan rules. Importantly, the plan distribution includes plans with immediate vesting, so our model accurately reflects the share of job spells that are not subject to vesting requirements.[37]

Finally, we calibrate participation, contribution, and separation rates to observed behavior. In particular, we estimate participation rates[38] and average contribution rates among employees in our balanced panel of plans during our sample period, splitting by decadal age buckets and income quintiles (with cut points again computed within age buckets and across plans). Using the same age and income splits, we estimate age-, income-, and tenure-specific separation hazard rates. We assume that the simulated worker participates, contributes, and separates according to these estimates, which we show in Figures B.2-B.3. The correlation between income and separation hazards in Figure B.3 is the key model input that generates income heterogeneity in forfeiture losses: because lower-income workers separate at higher rates, they hold more jobs during their careers, have shorter tenures

[36]We use the glide path (i.e., the age-dependent allocation shares to equity and bonds) employed in Vanguard's Target Retirement Funds.

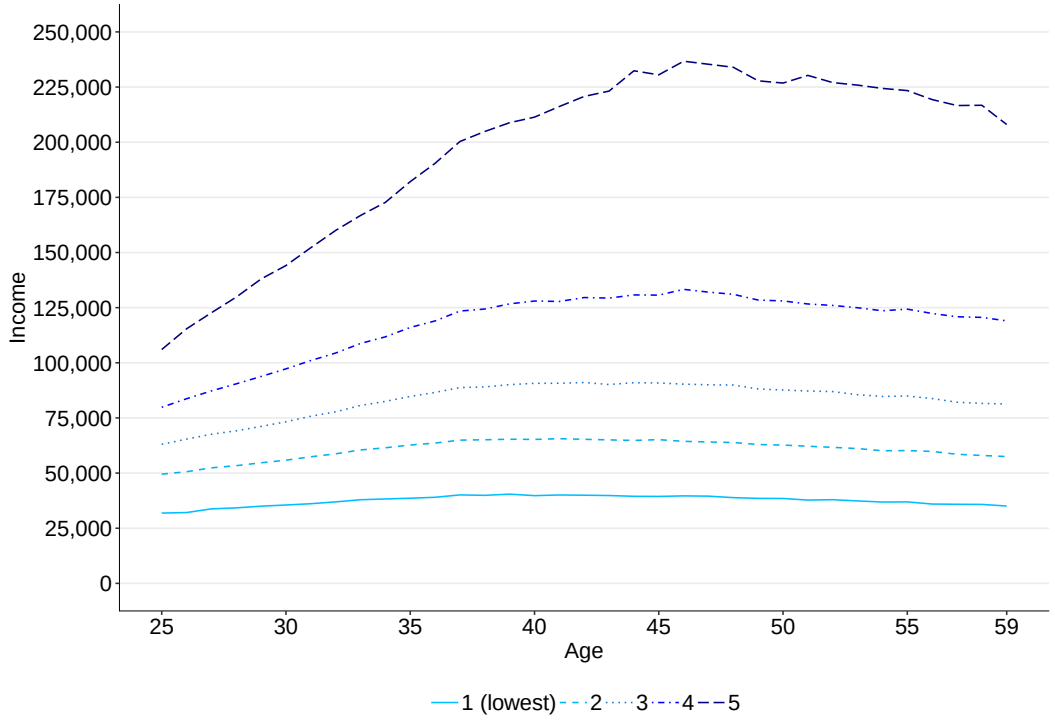
[37]This is a change from Section 4.2, where we drop immediate-vesting plans (about 25% of plans in our sample) in order to estimate the distributional effects of nontrivial vesting requirements.

[38]For the subset of plans in the nondiscrimination-testing data, we can observe non-participants and therefore calculate participation rates.

at separation, and are more likely to suffer forfeitures.

By using separation hazards that are uniform across plans, we rule out strategic separation behavior where quit rates increase at the end of a plan's vesting schedule. This assumption simplifies the model but also reflects our empirical findings: in Sections 5-6, we demonstrate that vesting requirements do not exert causal retention effects.

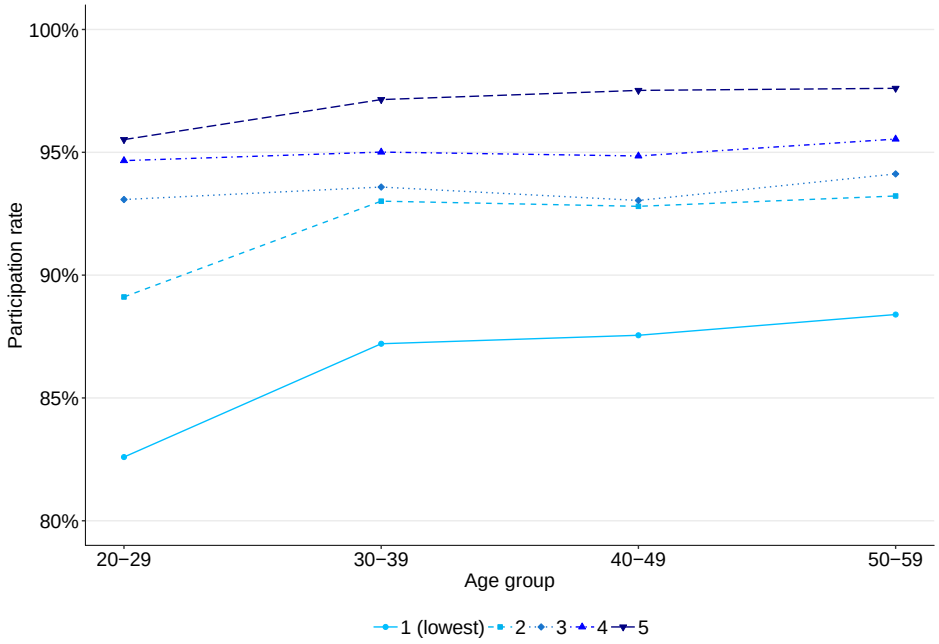
Figure B.1: Lifetime simulation calibration:
Quintile-specific lifetime income paths



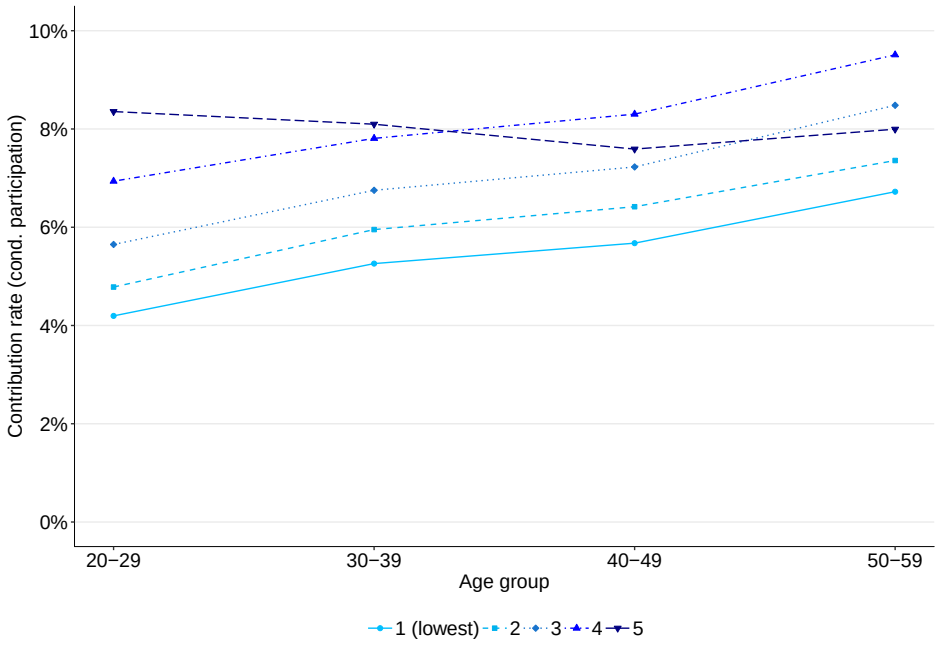
Note: This plots shows quintile-specific lifetime income paths, which we use to calibrate our lifetime simulation model in Section 4.3. We construct these paths using the 2022 cross-sectional income distribution among our balanced panel of plans. For each age from 25 to 59, we split employees of that age into income quintiles and compute average annual income within each quintile (to limit the influence of the right tail, we use median rather than average income for the top quintile).

Figure B.2: Lifetime simulation calibration:
Participation and contribution rates by age and income quintile

(a) Participation rates

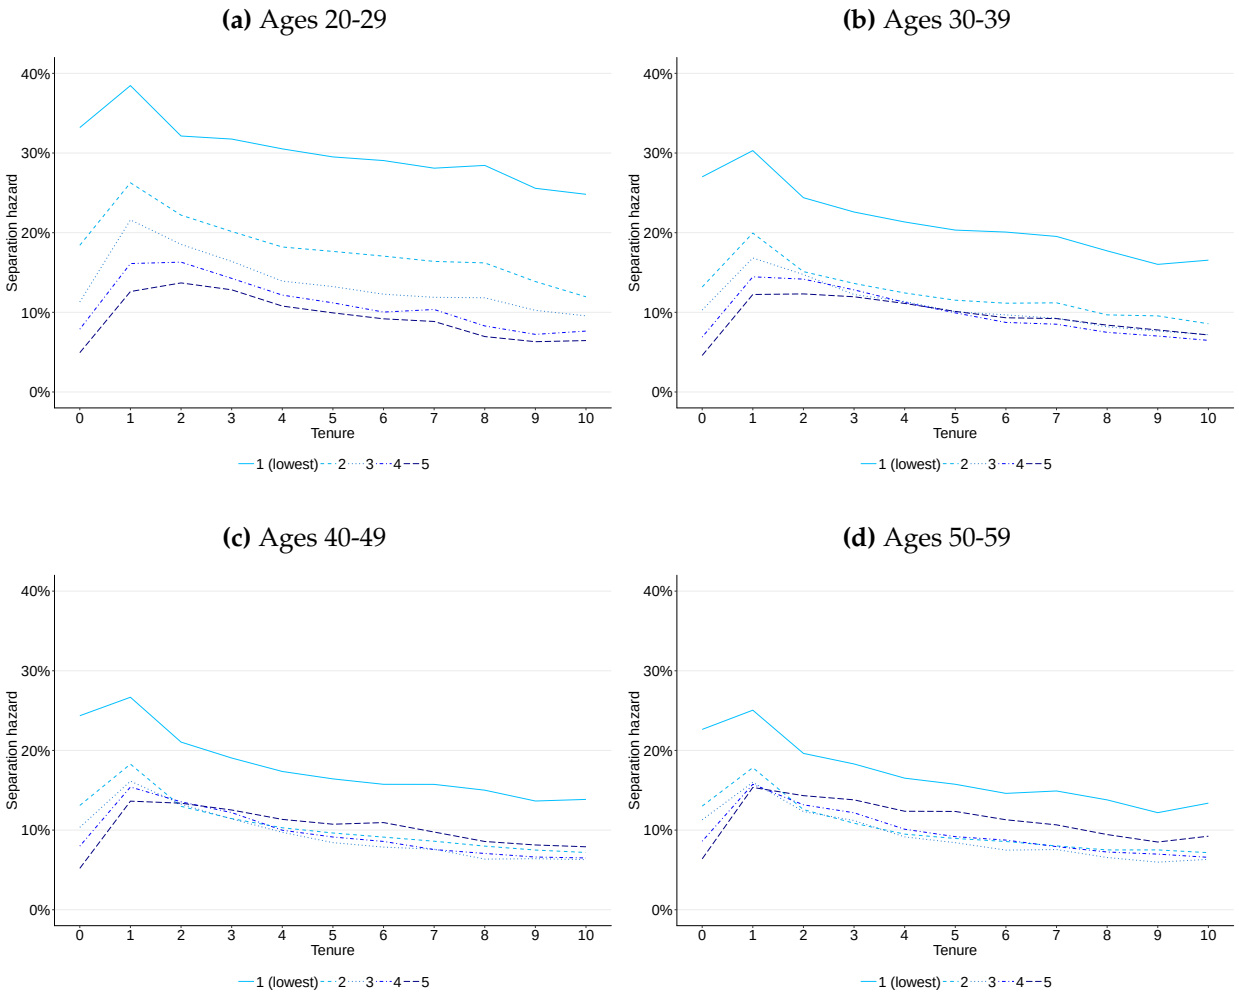


(b) Contribution rates



Note: These plots shows age- and income-specific participation and average contribution rates among employees in our balanced panel of plans, which we use to calibrate our lifetime simulation model in Section 4.3. The average contribution rates are conditional on participation. The top-quintile average contribution rate is decreasing in age because IRS limits on annual 401(k) contributions become increasingly likely to bind.

Figure B.3: Lifetime simulation calibration:
Separation hazard rates by age and income quintile



Note: This plot shows age-, income-, and tenure-specific separation hazard rates among employees in our balanced panel of plans, which we use to calibrate our lifetime simulation model in Section 4.3. Tenure is measured as whole years since hire (employees who have not yet reached their first hire anniversary have tenure of 0 years, employees who have reached their first hire anniversary but not their second have tenure of 1 year, etc.).

Appendix C Participant survey

To gauge workers' awareness and understanding of 401(k) vesting schedules, we use data from a Vanguard survey of defined-contribution plan participants. Because certain relevant data attributes are only available for plans using Vanguard's nondiscrimination-testing services (including W-2 income), the survey was restricted to the roughly two-thirds of plans that are in the nondiscrimination-testing data. The survey included questions assessing participants' general understanding of vesting schedules, their specific knowledge of vesting rules at their current plan, and their specific knowledge of non-vesting rules at their current plan. The survey also solicited basic demographic information and asked standard financial literacy questions. Each survey invitation was associated with a unique client identifier, allowing us to link survey responses with plan- and participant-level information in the administrative data.

The survey was conducted in October 2024 by sending email invitations to a random sample of 100,000 current participants. The sample receiving survey invitations was restricted to currently employed participants (i.e., it did not include participants who have separated from their employer but still have a positive account balance). Table C.1 gives summary statistics for the survey invitees and respondents. We disregard incomplete responses and classify participants as "respondents" only if they both started and finished the survey. In total, 1,018 participants finished the survey, implying a response rate of about 1%. We obtain age, tenure, and plan-level vesting information from the administrative recordkeeping data as of October 2024. We use 2023 income from the administrative data if available and respondents' self-reported income otherwise. The invitee and respondent samples are roughly balanced in terms of income and plan-level vesting rules, but older and longer-tenured participants were more likely to complete the survey and are overrepresented in the respondent sample.

Figure C.1 shows the email invitation and Table C.2 shows the survey questions.

Table C.1: Summary statistics for survey sample

	Random sample receiving survey invitation	Survey respondents
Number of participants	100,000	1,018
Number of plans	806	285
Age		
20-29	12%	2%
30-39	24%	8%
40-49	24%	17%
50-59	25%	37%
≥ 60	14%	35%
Tenure		
0-2	31%	23%
3-4	11%	10%
5-6	10%	8%
7-10	14%	12%
≥ 11	34%	46%
Income		
$<$ 30,000	5%	5%
30,000-59,999	22%	22%
60,000-89,999	25%	23%
90,000-119,999	16%	18%
$\geq$ 120,000	31%	33%
Vesting schedule		
Immediate	41%	44%
2-year cliff	4%	4%
3-year cliff	16%	19%
3-year graded	3%	3%
5-year graded	15%	11%
6-year graded	7%	7%
Other	15%	12%

Note: Respondents are those who started and finished the survey. We obtain age, tenure, and vesting information from the administrative recordkeeping data as of October 2024. If a participant's plan has more than one vesting schedule (which can occur if the employer makes both matching and nonmatching contributions), we associate the participant with the longest vesting schedule at the plan. We use annualized 2023 income from the administrative data if available and respondents' self-reports otherwise.

Figure C.1: Email invitation

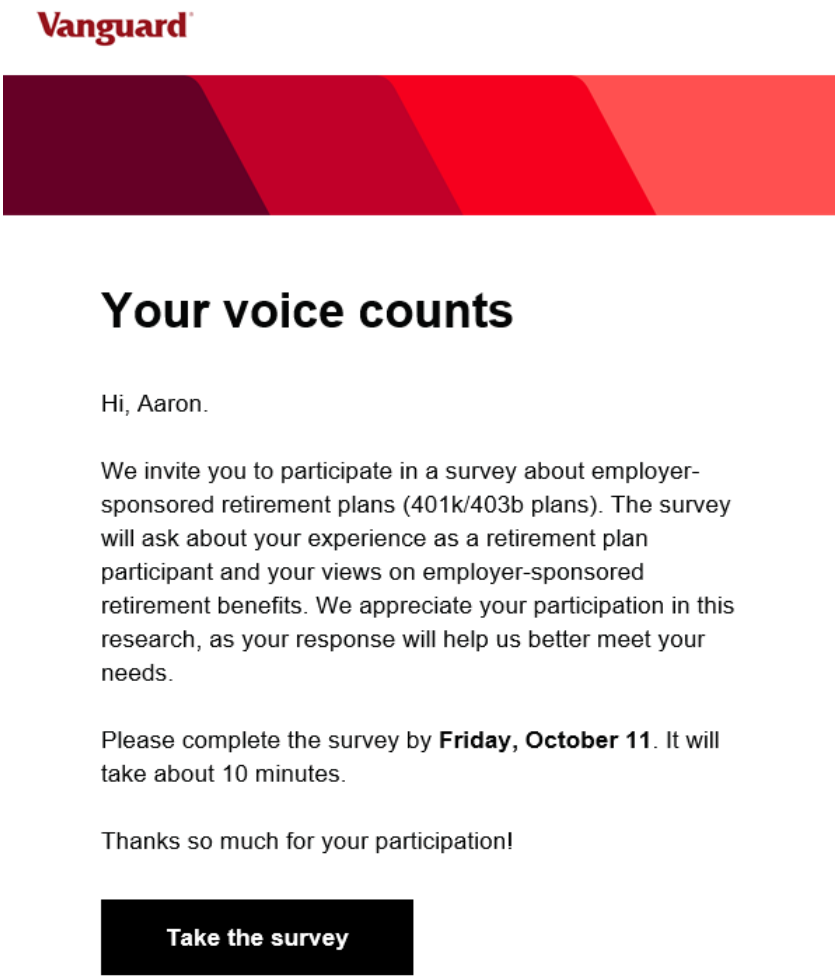


Table C.2: Survey questionnaire

Plan Feature Awareness	
Loans	Does your current plan allow you to take loans from your account? This means that your employer allows you to take a part of your retirement balance as a loan which you need to pay back over time. [Yes, No, I do not know/I am not sure]
Hardship Withdrawals	Does your current plan allow you to take hardship withdrawals from your account? This means that your employer allows you to take money from your retirement balance in case of a documented financial need. [Yes, No, I do not know/I am not sure]
Employer Contributions	Does your current employer make contributions to your plan? [Yes, No, I do not know/I am not sure]

Contribution Type	For this question, we refer to two different types of employer contributions. Matching contributions: An employer contributes some amount of money to your retirement account only if you also contribute. Non-matching contributions: An employer contributes some amount of money to your retirement account regardless of whether or not you contribute. Which type(s) of contributions does your employer make to your account? [Matching contributions only, Non-matching contributions only, Both matching and non-matching contributions, My employer does not make contributions into my plan, I do not know/I am not sure]

Vesting Understanding	
Vesting Definition	In the context of an employer-sponsored retirement plan like yours, what is a vesting schedule? [A rule specifying how long you must work at your employer to be eligible to participate in the plan, A rule specifying how long you must work at your employer to own the contributions your employer makes to your account, A rule specifying how long you must work at your employer to be able to withdraw money from your account without penalty, A rule that determines the amount of money your employer contributes to your account, None of the above, I do not know / I am not sure]

Preamble

A vesting schedule is a rule specifying how long you must work at your employer to own the contributions your employer makes to your account.

At plans with vesting schedules, employees who leave their employer before the end of the vesting schedule lose some or all of the contributions their employer made to their plans.

Your plan may or may not have a vesting schedule. In plans without a vesting schedule, you always own all of your employer contributions.

There are two types of vesting schedules: cliff schedules and graded schedules.

Cliff schedules: Employees own 0% of their employer contributions until a specific date, when their ownership increases to 100%.

For example: In a 3-year cliff schedule, when you reach 3 years of employment, you own all of your employer contributions. Before that, you own none of your employer contributions. The length of the schedule is 3 years.

Graded schedules: Employees' ownership of employer contributions increases step-by-step each year until reaching 100%.

For example: In a 4-year graded schedule, you might own 25% of your employer contributions after 1 year, 50% after 2 years, 75% after 3 years, and 100% after 4 years. The length of the schedule is 4 years.

Vignette 1

Please consider the following scenario.

Maria has been working at her employer for 3 years. She has participated in the 401(k) plan offered by her company and has saved $4,000 to date. Her employer has matched all of her contributions, so she has received $4,000 in employer contributions and has a total of $8,000 in her account.

Her plan has a 2-year graded vesting schedule. If she leaves her job before reaching 1 year of employment, she loses all of her employer contributions. If she leaves her job after 1 year but before 2 years of employment, she loses half (50%) of her employer contributions.

Let's suppose Maria leaves her job next month. What would be the total balance in her account after she leaves her job? [$3000, $4000, $5000, $6000, $7000, $8000, I do not know / I am not sure]

Vignette 2	Please consider the following scenario. Mark has been working at his employer for 2 years and 1 month. He has participated in the 401(k) plan offered by his company and has saved $2,000 to date. His employer has matched half of his contributions, so he has received $1,000 in employer contributions and has a total of $3,000 in his account. His plan has a 3-year cliff vesting schedule. If he leaves his job before reaching 3 years of employment, he loses all of his employer contributions. Let's suppose Mark leaves his job next month. What would be the total balance in his account after he leaves his job? [$500, $1000, $1500, $2000, $2500, $3000, I do not know / I am not sure]

Vesting Awareness	
Current Plan Vesting	Does your retirement plan have a vesting schedule for employer contributions? [Yes, No, I do not know/I am not sure]
Vesting Schedule Type	Is your plan's vesting schedule a cliff schedule or a graded schedule? [Cliff schedule, Graded schedule, I do not know/I am not sure]

Vesting Duration	How long is your plan's vesting schedule? In other words, how long do you need to work at your employer to own 100% of your employer contributions? [1 year, 2 years, 3 years, 4 years, 5 years, 6 years, 7 or more years, I do not know/I am not sure]
Vesting Status	Are you fully vested in all of your employer contributions? Being fully vested means that you own 100% of the contributions your employer has made to your account.[Yes, No, I do not know/I am not sure]

Demographics	
Income	What is your household's estimated annual pre-tax income this year? [Less than $30,000, $30,000 to $59,999, $60,000 to $89,999, $90,000 to $119,999, $120,000 to $149,999, $150,000 to $179,999, $180,000 or more, Prefer not to say]
Race/Ethnicity	What is your race/ethnicity? [White, Hispanic, Black or African American, American Indian or Alaska Native, Asian, Native Hawaiian or Pacific Islander, Other, Prefer not to say]
Age	What is your age group? [18-24 years, 25-34 years, 35-44 years, 45-54 years, 55-64 years, 65 years or older, Prefer not to say]

Gender	How do you identify? [Male, Female, Other, Prefer not to say]
Education	What is the highest level of schooling you have completed, or the highest degree you have received? [Less than high school, High school graduate, 2 year degree, 4 year degree, Graduate degree, Prefer not to say]
Financial Literacy	
Inflation Concept	Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After one year, how much would you be able to buy with the money in this account? [More than today, Exactly the same as today, Less than today, I do not know / I am not sure]
Compound Interest	Suppose you had $100 in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow? [More than $102, Exactly $102, Less than $102, I do not know / I am not sure]
Stock Diversification	Please say whether this statement is true or false: Buying a single company's stock usually provides a safer return than a stock mutual fund. [True, False, I do not know/I am not sure]

Bond Prices	If interest rates fall, what would happen to the prices of bonds? [They would rise, They would fall, They would stay the same, None of the above, I do not know / I am not sure]
401(k) Taxation	Thinking about saving in a traditional 401(k) account and withdrawing after age 60, which of the following is true: [You pay income tax on what you save now but do not pay taxes on withdrawals, You pay no income tax on what you save now but pay income taxes on withdrawals, You pay income tax both on what you save now and on withdrawals, You pay no income tax in what you save now or on withdrawals, I do not know / I am not sure]