Impatience or illiquidity? The case of early 401(k) withdrawals*

Aaron Goodman
Vanguard

March 2025

Abstract

Over a third of 401(k) participants cash out their balances when they leave their jobs. These pre-retirement withdrawals could arise from behavioral biases or liquidity constraints, but empirical evidence distinguishing between alternative explanations is lacking. Using a novel paycheck-level income panel covering millions of 401(k) participants, I show that workers who experience greater sub-annual income volatility are significantly more likely to cash out their balances at separation. The relationship between earnings volatility and cashout probability is conditional on annual income, driven primarily by negative earnings shocks, and most pronounced among low- and middle-income workers. The empirical results are consistent with a model in which workers use cashouts as a consumption-smoothing measure when income shocks deplete their liquid resources over the course of an employment spell.

*I thank John Beshears, Guillermo Carranza, Marcos Catao, James Choi, Taha Choukhmane, Jeff Clark, Paulo Costa, Peter Ganong, Fiona Greig, Kelly Hahn, David Laibson, Pascal Noel, Cormac O'Dea, David Pakula, Lawrence Schmidt, Fu Tan, Nicky Zhang, and colleagues at Vanguard's Investment Strategy Group for valuable feedback. The views expressed here are those of the author 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

A significant share of assets in defined-contribution retirement accounts are withdrawn before retirement. The implications of this fact – for public policy, for the design of employer-sponsored retirement plans, and for analysis of household saving behavior more broadly – depend crucially on the reasons for early withdrawals. If retirement savings leakage arises from behavioral biases (e.g., present bias), then early-withdrawal penalties raise welfare. Alternatively, if households tap retirement accounts because they lack the liquid resources to meet current consumption needs, then they benefit from policies that facilitate liquid wealth formation and make retirement assets available during emergencies.

This paper presents the strongest evidence to date that liquidity constraints are an important explanation for early withdrawals. The empirical analysis centers on a novel paycheck-level income panel covering 3 million 401(k) plan participants across a 15-year period. Because I construct the paycheck panel from administrative 401(k) recordkeeping data, I can relate granular features of participants' income histories to their saving and withdrawal decisions. When I do so, a consistent finding emerges: participants with greater sub-annual income volatility are significantly more likely to take penalized early withdrawals from 401(k) accounts. The relationship between earnings volatility and withdrawal rates is conditional on annual income, driven primarily by negative earnings shocks, and most pronounced among low- and middle-income participants. These empirical results are consistent with a conceptual framework in which workers take early 401(k) withdrawals when income shocks create liquidity shortages.

To formalize the theoretical framework and discipline the empirical analysis, I develop a model of retirement saving and withdrawal decisions that reflects the institutional reality of the U.S. defined-contribution (DC) system. In the model, a worker earns stochastic labor income over multiple periods and saves for retirement in a 401(k) account. The worker may withdraw from the 401(k) account before retirement, but early withdrawals incur a penalty and are only permitted in a designated subset of periods during working life. The designated early-withdrawal periods represent times when the worker changes jobs and thus becomes eligible to take a penalized cash distribution from her 401(k) account. Because withdrawals during employment are subject to relatively strict regulation, the post-separation cashouts on which the model focuses are the most common form of 401(k) leakage: in 2023, 33% of all separators in Vanguard-administered plans cashed out (Clark, 2024). The modeled worker also has access to a lower-yielding liquid savings account,

from which withdrawals can be taken at any time without penalty.

The model makes three core predictions that structure the empirical analysis and shape my interpretation of the empirical results. First, negative income shocks cause cashouts. When the worker receives sufficiently poor draws from the stochastic income process, the consumption-smoothing benefit of cashing out becomes larger than the early-withdrawal penalty. Second, mean-zero income volatility causes cashouts. Even when negative income shocks are accompanied by equally sized positive shocks during the same employment spell, the worker uses her liquid wealth to smooth consumption, has less liquidity at separation, and may cash out in order to replenish the liquid account before beginning her next job. Third, the relationship between income volatility and cashout propensity holds in a baseline case where the worker is a standard exponential discounter, but becomes stronger when the worker is present-biased (Laibson, 1997). This is because liquidity is partly endogenous: present-biased workers build smaller precautionary buffers in the liquid account, have less liquidity on hand when income shocks arrive, and must resort to penalized 401(k) withdrawals more often.

Motivated by the model of withdrawal behavior, I build a dataset that is uniquely suited to test its predictions. High-frquency income observations are crucial: even short-lived negative earnings shocks can deplete liquidity buffers, and mean-zero volatility within a given year would be completely undetectable in annual income data. Using administrative 401(k) recordkeeping data from Vanguard, I rely on the fact that for most of the plans it administers, Vanguard tracks both the dollar amounts of employee contribution transactions and employees' underlying contribution rates. Combining the information on contribution amounts and rates allows me to infer individual paychecks. For example, a participant who made a $50 contribution with a 5% contribution rate on January 1 likely received a $1,000 paycheck on January 1. Focusing on cases where I can verify that the annual sum of inferred paychecks matches official W-2 income figures, I obtain complete paycheck-level income histories for 9.3 million particpant-years across 1,224 plans during the 2008-2022 period.

The paycheck data confirm the model's theoretical predictions. I first divide the sample of participants into those who are likely hourly workers and those who are likely salaried, based on the frequency with which their paycheck amounts change by at least $1 relative to the lagged paycheck amount.[1] Among those who separate from employment during the

[1] I thank Peter Ganong, Pascal Noel, and their coauthors in Ganong et al. (2024b) for providing statistics from external payroll data that help make this hourly-salaried classification as accurate as possible.

sample period, low- and middle-income workers in the hourly group have cashout rates about 12 percentage points higher than salaried workers with similar annual incomes. To measure volatility quantitatively, I work at the monthly level and compute the standard deviation of monthly income *within* a given participant-year. In regressions that control for annual income, age, employment tenure, and plan-level fixed effects, I find that a 10% increase in the standard deviation of monthly income raises post-separation cashout probability by 3.5 percentage points for middle-income workers (a 10% increase from the baseline cashout rate of 35%). When I control separately for the second moment and the third moment (i.e., skewness) of monthly income, I find that asymmetric negative shocks raise cashout rates by about twice as much as mean-zero volatility.

Finally, because income volatility is generally not randomly assigned, I offer two additional pieces of evidence supporting the model's causal liquidity mechanism. First, I identify cases where many participants in the same plan separate in the same month, likely indicating that the separations resulted from firm-wide layoffs (Jacobson et al., 1993). Since unemployment insurance provides liquidity to laid-off workers and 401(k) withdrawals mechanically reduce UI benefits under many states' UI rules,[2] workers motivated by liquidity considerations should be less likely to cash out after being laid off. I find that participants who separate in apparent plan-wide layoff events are in fact significantly less likely to cash out (controlling for age, income, tenure, plan fixed effects, and realized income volatility during the employment spell). The negative layoff effect is larger for lower-income workers (who have higher UI replacement rates) and during the 2020-2021 period (when pandemic-era supplements made UI benefits more generous). Second, I adopt an instrumental-variable strategy, using negative income shocks among coworkers as a labor-demand-driven instrument for own income shocks (Koustas, 2018; Ganong et al., 2023). The IV estimates are noisier but similar in magnitude to the baseline estimates.

I contribute to an empirical literature studying pre-retirement leakage from retirement accounts. Most prior work in this literature computes leakage rates using federal tax-return data.[3] Goodman et al. (2021) find that from 2003 to 2015, 22% of contributions to DC and Individual Retirement Arrangement (IRA) accounts by working-age individuals leaked out as pre-retirement withdrawals.[4] Argento et al. (2015) show that each year during the

[2] See the discussion in Section 5.1.

[3] Choi et al. (2024) is an exception: using recordkeeping data covering nine firms between 2005 and 2011, they find that 42% of 401(k) balances are cashed out upon separation.

[4] I share a last name with the first-listed author on this paper but we are not related. I was not invited to collaborate on Goodman et al. (2015).

2004-2010 period, roughly 8% of households with DC accounts took penalized distributions. Choukhmane et al. (2024) document that early withdrawals are significantly more common among minorities and lower-income workers and contribute to racial and income gaps in DC wealth accumulation. Amromin and Smith (2003) compute leakage rates using less recent data from the 1990s.

In addition to studying average leakage rates, Goodman et al. (2021), Argento et al. (2015), and Amromin and Smith (2003) each find evidence that households are more likely to take early withdrawals in years where their annual income drops by a large amount. By using recordkeeping data to study the effect of income volatility, I provide a more thorough analysis that overcomes several fundamental limitations of annual tax-return data. First, the correlation between income shocks and withdrawals in tax-return data is partly mechanical: job changes make workers eligible to take 401(k) distributions and may simultaneously reduce annual income if there is a transition or unemployment period between jobs. I *condition* on separation, then relate cashout decisions to income volatility *during* the preceding employment spell. Second, the sub-annual volatility I study is fundamentally different (and likely more common) than large shocks to annual income. I can detect the short-lived negative earnings shocks and mean-zero, within-year volatility that matter for workers' liquidity positions but are not observable in annual income data. Finally, 401(k) recordkeeping data alleviates some aggregation issues in tax-return data, like the fact that distributions from defined-contribution and defined-benefit accounts are not separately reported on Form 1099-R and must be inferred using other information on the return (Goodman et al. 2021 discuss this issue in detail).

I show that liquidity constraints are an important explanation for early withdrawals, but many other factors may cause leakage, and behavioral biases are likely among them. Policymakers setting withdrawal rules therefore face a tradeoff between mitigating behavioral biases and providing flexibility to households with high marginal utility of consumption. Beshears et al. (2025) study this tradeoff in a generalized planning problem (building on the modeling framework of Amador et al. 2006). They find that the optimal retirement savings system is well approximated by a two-account structure: one account is completely illiquid and mitigates present bias by forcing a high level of mandatory savings, while a second liquid account provides limited flexibility to patient households with marginal-utility shocks. The optimality of mandating high illiquid savings levels depends in part on an assumption that present bias is sufficiently heterogeneous relative to the distribution of marginal-utility shocks. The substantial monthly income volatility I document among

401(k) participants provides a micro-foundation for the Beshears et al. (2025) marginal-utility shocks and shows that these shocks have wide empirical variance.[5] An important avenue for future work is to estimate precisely the empirical distribution of present bias,[6] allowing for data-driven evaluations of the commitment-versus-flexibility tradeoff.

I also contribute to a literature studying the optimal design of employer-sponsored 401(k) plans within the DC system. Since the seminal contribution of Madrian and Shea (2001), a substantial body of empirical work has shown that automatic enrollment increases participation and savings rates in 401(k) plans (Choi et al., 2002, 2004; Beshears et al., 2009). More recent studies show that the beneficial effects of automatic enrollment partially dissipate over time as workers reduce non-retirement saving (Choukhmane and Palmer, 2024), take on debt (Beshears et al., 2024a), and cash out at separation (Choi et al., 2024). The empirical link between income volatility and cashout probability suggests liquidity constraints are a partial explanation for attenuated long-run savings effects. Theoretically, the model I develop demonstrates the value of plan-design choices that account for workers' liquidity needs: the same behavioral biases that inhibit retirement saving also cause workers to undersave in precautionary liquid accounts (Ameriks et al., 2007; Ganong and Noel, 2019). Promising interventions include those that enroll workers in employer-sponsored short-run savings accounts (Berk et al., 2023) and allow limited amounts of emergency 401(k) liquidity during workers' employment spells (Beshears et al., 2024b).[7]

Finally, I contribute to a growing literature on heterogeneous labor income risk (Guvenen et al., 2014, 2021; Busch et al., 2022; Meeuwis et al., 2025). Whereas this literature relies primarily on annual income data and studies earnings risk that arises largely from job changes and unemployment spells, I show that sub-annual volatility within a given employment spell can have important household financial implications. In using paycheck-level data and analyzing monthly earnings volatility, I follow Ganong et al. (2024b), who are the first to do so with U.S. administrative data. Relative to the payroll-processing data in Ganong et al. (2024b), the earnings information I construct from 401(k) recordkeeping

[5]Credit card transaction data (Farrell and Greig, 2015; Farrell et al., 2019) and household survey data (Fulford and Low, 2024) suggest expense shocks are an important determinant of liquidity constraints and that monthly expenses are even more volatile than monthly income. The income volatility I document could therefore understate the total variance of marginal-utility shocks.

[6]See the literature reviews in Lockwood (2020) and Cohen et al. (2020). Of primary importance is research that estimates the distribution of present bias within a given sample using a consistent methodology (Paserman 2008 and Meier and Sprenger 2015 are examples), rather than relying on cross-study variance in average estimates.

[7]See the discussion of 401(k) liquidity rules in Section 6.

data is valuable because it allows me to relate income volatility to workers' saving and withdrawal decisions. To my knowledge, this paper is the first to link paycheck-level income data in the U.S. with worker-level financial behaviors.

The remainder of the paper is organized as follows. Section 2 develops the model of retirement saving and withdrawal decisions. Section 3 describes the administrative 401(k) recordkeeping data and the process by which I infer paychecks. Section 4 presents the main empirical results linking sub-annual income volatility to early 401(k) withdrawals. Section 5 uses layoffs and instrumental-variable analysis to provide additional evidence in support of a causal liquidity mechanism. Section 6 discusses policy implications and Section 7 concludes.

2 Model

2.1 Setup

To motivate the empirical analysis, I introduce a simple model of retirement saving and withdrawal decisions. Taking important features of the U.S. retirement system as given, the model describes the partial-equilibrium behavior of a worker who saves in a 401(k) account and seeks to maximize discounted lifetime utility. As in prior literature, I allow the worker to be naively present-biased (Laibson, 1997; Beshears et al., 2022). The distinguishing feature of the model is labor income volatility, which can shift the marginal utility of consumption during working life and potentially induce early withdrawals from the 401(k) account.[8]

The worker's life consists of T periods. In each period t, the worker's objective is to maximize discounted utility, which is given by

$$u(c_t) + \beta \sum_{x=t+1}^{T} \delta^{x-t} u(c_x) \tag{1}$$

The exponential discount factor is $\delta < 1$ and $\beta \in (0, 1]$ controls the degree of present bias; the benchmark case of $\beta = 1$ represents a worker who is not present-biased.

[8] In similar dynamic settings, Beshears et al. (2025) and Amador et al. (2006) allow for time-varying taste shocks θ that shift marginal utility. The random income shocks here can be seen as a micro-foundation for such taste shocks (though with a different functional form: Beshears et al. 2025 and Amador et al. 2006 consider multiplicative shocks $\theta u(c)$, while I consider shocks to the consumption argument c). Moser and de Souza e Silva (2019) consider persistent heterogeneity in earnings ability but not time-varying, within-person shocks.

For $t < t_R$, the worker is employed and earns stochastic labor income y_t that is independently distributed across periods. For $t \geq t_R$, the worker is retired and receives fixed Social Security income y_R. Consumption may differ from current income because the worker has access to two savings accounts. The first account, with balance at the beginning of period t denoted by a_t^L, is a completely liquid account from which the worker can withdraw at any time without penalty. The only constraint applying to the liquid account is that its balance may not fall below a fixed floor, $a_t^L \geq \underline{a^L}$. I consider a simple no-borrowing case, so $\underline{a^L} = 0$.[9] The return on the liquid account is normalized to zero.

The second savings account is a partially illiquid 401(k) account, with balance at the beginning of period t denoted by a_t^I. The 401(k) account must always have a weakly positive balance and earns a return of r between periods. The worker can only contribute to the account during working life for $t < t_R$. Contributions to the account are matched by the worker's employer at rate m up to a per-period cap of $\bar{m}$. During the retirement period for $t \geq t_R$, the worker may withdraw from the account without penalty. Additionally, for a subset of periods during working life $t \in W$, the worker may withdraw from the 401(k) account but must pay a proportional early-withdrawal penalty of τ.[10] The pre-retirement withdrawal periods W capture instances where the worker changes jobs and thus becomes eligible to take a penalized cash distribution.[11] Limiting early withdrawals to periods with job changes reflects the institutional reality of the DC system, where all workers are eligible to take withdrawals upon separation but in-service withdrawals are more restricted and much less common.[12]

Let s_t^I be the worker's 401(k) saving at time t, with positive values indicating contributions and negative values indicating withdrawals. Additionally, let $V_t(a^L, a^I)$ be the period-t value function representing the value of optimal behavior from period t onward, taking the two state variables (liquid and illiquid account balances) as arguments. Because

[9] A straightforward extension allowing borrowing would have $\underline{a^L} < 0$ and negative balances growing in magnitude from t to $t+1$ because of interest costs. The model's main qualitative conclusion would not change: as long as the interest rate paid on debt exceeds the interest rate earned on 401(k) savings (the relevant empirical case if short-term borrowing occurs on credit cards), negative income shocks would still induce early 401(k) withdrawals (either to avoid credit card borrowing or to pay off a costly stock of existing credit card debt).

[10] I do not formally model taxes or the tax-deferral benefits of saving in a 401(k) account. The early-withdrawal penalty τ could be interpreted as representing both the nominal withdrawal penalty and the cost of forgoing tax-deferred investment returns.

[11] A simplifying assumption of the model is that these job changes occur with certainty in predetermined periods rather than stochastically.

[12] In the empirical sample, 34.6% of separators take a cash distribution but in-service hardship withdrawals occur in only 1.7% of participant-years. See the discussion in Section 4.1.

any present bias with $\beta < 1$ is naive, the value function is computed assuming that the worker will make optimal choices according to $\beta = 1$ from period t onward. The fully stated problem for the worker at any time t is thus:

$$\max_{c_t, s_t^L, s_t^I} \; u(c_t) + \beta\delta V_{t+1}(a_{t+1}^L, a_{t+1}^I), \tag{2}$$

with state evolution given by:

$$c_t + s_t^L + s_t^I + p_t = y_t, \tag{3}$$

$$a_{t+1}^L = a_t^L + s_t^L, \tag{4}$$

$$a_{t+1}^I = (1+r)(a_t^I + s_t^I + \max(0, m * \min(\bar{m}, s_t^I))), \tag{5}$$

and constraints given by:

$$s_t^L \geq -a_t^L, \tag{6}$$

$$s_t^I \begin{cases} \in [-a_t^I, 0] & \text{if } t \geq t_r \\ \geq -a_t^I & \text{if } t \in W \\ \geq 0 & \text{otherwise,} \end{cases} \tag{7}$$

$$p_t = \begin{cases} \tau s_t^I & \text{if } s_t^I < 0 \text{ and } t \in W \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

The worker chooses consumption c_t and 401(k) saving s_t^I; these two choices, along with the budget constraint in (3), imply liquid saving s_t^L. The 401(k) account earns a matching contribution if $s_t^I > 0$ and grows with net return r. The constraints in (6)-(8) require that neither account balance fall below zero, that 401(k) withdrawals occur only during retirement or during the designated pre-retirement periods W, and that the worker pay a penalty for any pre-retirement 401(k) withdrawals.

2.2 Simple three-period case

To build intuition for the model's key results, I first consider the simplest meaningful case: two working periods, one retirement period, and one job change. Formally, this corresponds to $T = 3$, $t_R = 3$, and $W = \{2\}$. At $t = 1$, the worker earns labor income from

distribution $F_1(y)$. At $t = 2$, the worker changes jobs, earns labor income from distribution $F_2(y)$, and is allowed to cash out the 401(k) balance accumulated in the first job. Finally, at $t = 3$, the retired worker consumes Social Security income y_R and all remaining wealth.

2.2.1 Cashout decision at $t = 2$

Proceeding by backward induction, consider the worker's behavior at $t = 2$. Let a_2^L and a_2^I be the account balances entering $t = 2$ and y_2 be the realized income draw. Because 401(k) withdrawals are not penalized at $t = 3$, there is no reason to forgo the 401(k) return r by carrying a liquid account balance into $t = 3$. The worker therefore depletes the liquid account and the budget constraint in (3) becomes $c_2 = y_2 + a_2^L - s_2^I - p_2$. A necessary[13] condition for the worker to choose $s_2^I < 0$ and take a penalized 401(k) withdrawal is:

$$u'\left(y_2 + a_2^L\right) > \beta\delta\left(\frac{1+r}{1-\tau}\right) u'\left(y_R + (1+r)a_2^I\right). \tag{9}$$

The left-hand side of (9) gives the marginal benefit of reducing s_2^I from zero to a small negative value. The right-hand side gives the corresponding marginal cost, which must be strictly smaller to justify a cashout. Inspecting (9), we can see that three key factors govern the cashout decision.

- *Current liquidity*: When liquid resources $y_2 + a_2^L$ are low, the marginal utility of consumption at $t = 2$ is high, incentivizing the worker to pull consumption forward in time by withdrawing from the 401(k) account. Low liquidity can arise from an exogenous shock to labor income y_2. It can also arise endogenously through a_2^L if the worker chose a low level of liquid saving at $t = 1$.

- *Future liquidity*: The worker's illiquid retirement wealth (which will become liquid at $t = 3$) is composed of Social Security income and the future value of the 401(k) balance, $y_R + (1 + r)a_2^I$. Mirroring the logic from the prior bullet, when the stock of illiquid wealth is large, the worker anticipates relatively low marginal utility of consumption at $t = 3$ and has a stronger incentive to tap 401(k) wealth before retirement. When y_R replaces a larger share of the worker's labor income, the worker therefore has a stronger cashout incentive.

[13] By taking a 401(k) withdrawal (and therefore not contributing to the 401(k) account) at $t = 2$, the worker forgoes the employer match. This makes possible a case where (9) holds but the employer match is sufficiently generous to make a cashout suboptimal. As the match rate m or match cap $\bar{m}$ approach zero, (9) becomes both necessary and sufficient for a cashout to occur.

- *Discount rate*: Withdrawn 401(k) funds forgo the market return r and incur the early-withdrawal penalty τ; both of these considerations raise the marginal cost on the right-hand side of (9). Fixing r, τ, and the exponential discount factor δ, the effective discount rate is determined by the present-bias parameter: as β decreases below 1, the cashout condition becomes more likely to hold.

The three cashout drivers (low current liquidity, high future liquidity, present bias) are "complements" in the sense in that each increases the marginal importance of the others. For example, consider a worker at $t = 2$ who does not cash out with some income draw $\hat{y}_2$. If the income draw is instead some smaller value $\tilde{y}_2 < \hat{y}_2$, the left-hand side of (9) increases. This reduction in current liquidity is more likely to satisfy the cashout inequality if a large y_R value or a small β value makes the right-hand side of (9) smaller to begin with. Similarly, a present-biased worker whose discount rate makes her predisposed to cashing out at $t = 2$ is even more likely to do so if her realized income draw y_2 is lower.

2.2.2 Numerical solutions

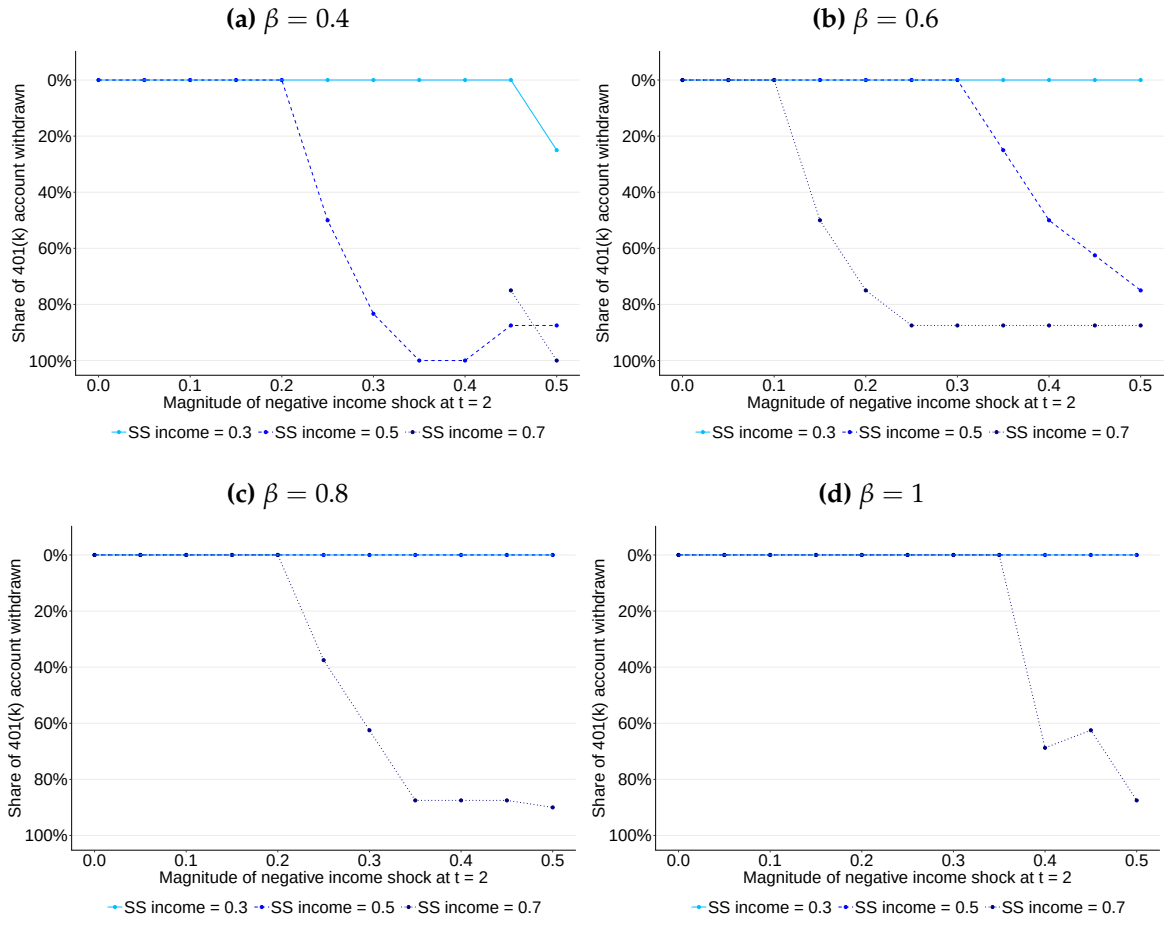
To make the theoretical considerations in the preceding section more concrete, I solve the three-period model numerically for an illustrative set of parameter values. To sharpen the intuition around the effect of income volatility, I use a simple binary income process. At $t = 1$, income is $y_1 = 1$ with certainty. At $t = 2$, income takes one of two values, occurring with equal probability: either $y_2 = 1 + \Delta$ or $y_2 = 1 - \Delta$ for some shock $\Delta \geq 0$.[14]

Figure 1 shows the share of the 401(k) balance a_2^I that the worker withdraws at $t = 2$, conditional on the negative income shock $y_2 = 1 - \Delta$ being realized. Each panel shows behavior under a different value for the present-bias parameter ($\beta \in \{0.4, 0.6, 0.8, 1\}$) and lines within each panel represent a different value for Social Security income ($y_R \in \{0.3, 0.5, 0.7\}$). Since $\mathbb{E}[y_1] = \mathbb{E}[y_2] = 1$, the y_R values can be interpreted as Social Security replacement rates, expressed as a fraction of expected lifetime labor income.

Consistent with the intuition from (9), cashouts in Figure 1 are more likely when i) a larger negative income shock reduces current liquidity, ii) higher Social Security income raises future liquidity, and iii) a smaller present-bias parameter decreases the worker's discount rate. Looking within each panel and fixing the degree of present bias shows the

[14]Other fixed parameters are as follows: utility is CRRA with coefficient of relatve risk aversion equal to 2, the exponential discount factor is $\delta = 0.97$, the 401(k) return is $r = 0.08$, the early withdrawal penalty is $\tau = 0.1$, the match rate is $m = 0.5$, and the match cap is $\bar{m} = 0.06$. Both the liquid account balance a_1^L and the 401(k) account balance a_1^I are initialized at zero.

Figure 1: Cashout choices at $t = 2$ in the three-period case, conditional on the negative income shock

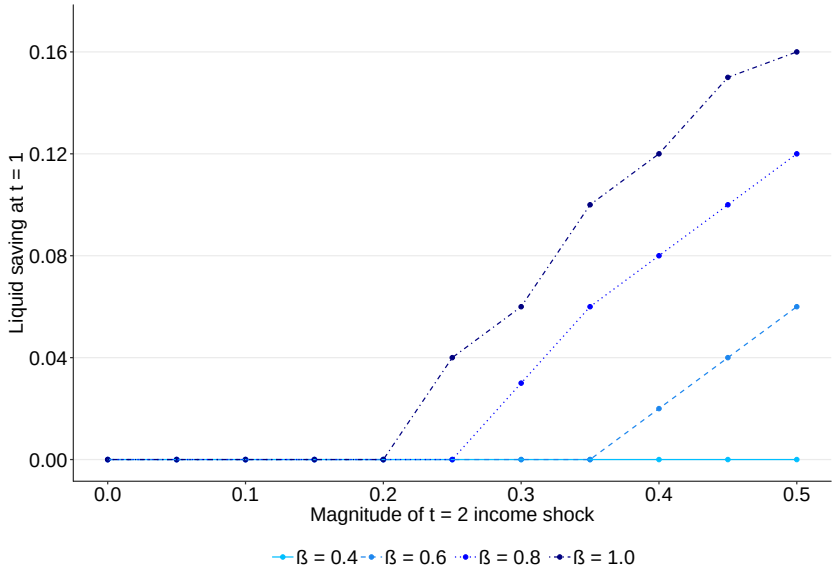


Note: This figure shows cashout choices made by the worker at $t = 2$ when the negative income shock $y_2 = 1 - \Delta$ realizes. The x-axis shows the magnitude of the Δ shock and the y-axis shows the share of the 401(k) account balance a_2^I that the worker chooses to take as a penalized withdrawal. If the worker does not take a withdrawal and $s_2^I \geq 0$, the plotted value is 0. Missing values indicate cases where the worker does not contribute to the 401(k) account at $t = 1$ and $a_2^I = 0$. Fixed model parameters are as follows: the coefficient of relative risk aversion is 2, the exponential discount factor is $\delta = 0.97$, the 401(k) return is $r = 0.08$, the early withdrawal penalty is $\tau = 0.1$, the match rate is $m = 0.5$, the match cap is $\bar{m} = 0.06$, and first-period income is $y_1 = 1$. The worker's saving choice set is discretized with steps of 0.01 between grid points. See Sections 2.1-2.2 for a description of the model and the simplified three-period case.

partial effects of current and future liquidity. In the benchmark $\beta = 1$ case, the worker only cashes out under the highest Social Security value $y_R = 0.7$ and for income shocks larger than 0.4. As β decreases from panels (d) to (a), the size of the income shock necessary to induce a cashout (for any given Social Security value) also decreases. At the opposite extreme for $\beta = 0.4$, large income shocks prompt 401(k) withdrawals even under the lowest Social Security value $y_R = 0.3$.

An important feature of the model is that the worker's liquidity position at $t = 2$ is not

Figure 2: Liquid saving at $t = 1$ in the three-period case



Note: This figure shows liquid saving choices made by the worker at $t = 1$. The x-axis shows the magnitude of the symmetric $t = 2$ income shock Δ, with $y_2 = 1 + \Delta$ and $y_2 = 1 - \Delta$ occurring with equal probability. The y-axis shows the worker's contribution to the liquid account at $t = 1$. Fixed model parameters are as follows: the coefficient of relatve risk aversion is 2, the exponential discount factor is $\delta = 0.97$, the 401(k) return is $r = 0.08$, the early withdrawal penalty is $\tau = 0.1$, the match rate is $m = 0.5$, the match cap is $\bar{m} = 0.06$, first-period income is $y_1 = 1$, and Social Security income is $y_R = 0.5$. The worker's saving choice set is discretized with steps of 0.01 between grid points. See Sections 2.1-2.2 for a description of the model and the simplified three-period case.

purely exogenous. Labor income y_2 is a random draw, but the liquid account balance a_2^L is chosen by the worker at $t = 1$. Figure 2 shows the endogenous accumulation of liquid savings at $t = 1$ (fixing Social Security income at the intermediate value of $y_R = 0.5$). For each β value, liquid saving rises with the size of the symmetric $t = 2$ income shock as the worker perceives a greater precautionary value of liquid wealth. The relationship between income volatility and liquid saving attenuates as β decreases; the most impatient worker with $\beta = 0.4$ never saves in the liquid account.

Two conclusions are worth drawing from Figure 2. First, within the model, cashouts are more likely for lower β values both because the effective discount rate at $t = 2$ is lower *and* because liquid savings a_2^L are lower. Second, when moving to empirical and policy analysis, cashout behavior will be partially determined by liquid account balances that are not observable in 401(k) recordkeeping data.

2.3 Expanded case with many periods

2.3.1 Setup and cashout decision

The simple three-period case is sufficient to establish the essential intuition around cashout decisions. Even so, it is worth briefly considering an expanded case that aligns more closely with the empirical setting of Section 4. The empirical analysis will consider mid-career 401(k) participants who earn income over many sub-annual periods (months), then separate from employment and decide whether to cash out. Participants' income histories are no longer observable in the Vanguard data following separation, but after making their cashout decision they presumably begin another job and eventually retire.

In the expanded case, I set $T = 15$, $t_R = 11$, and $W = \{5\}$. The periods preceding the cashout decision $t \leq 5$ represent the observable income history during a worker's employment spell at a Vanguard-administered plan. The worker separates and makes a cashout decision at $t = 5$, earns income at a different job for $6 \leq t \leq 10$, then retires at $t = 11$. I make retirement five periods long so that the relative lengths of working life and retirement are the same as in the three-period case. Finally, I let labor income y_t be independent across periods and follow a lognormal distribution with mean zero and standard deviation σ. I consider a range of σ values to illustrate comparative statics with respect to income volatility.[15]

The cashout decision at $t = 5$ is driven by the same factors from the three-period case (current liquidity, future retirement liquidity, present bias) but has two important differences. First, the worker's current liquidity position at $t = 5$ represents the cumulative effects of many random income draws during the first job. Second, because working periods remain before retirement, cashouts need not fund consumption exclusively in the current $t = 5$ period. If the worker has a low liquid account balance a_5^L and finds the continuation value of marginal liquid wealth to be high, she may cash out some of her 401(k) balance to deposit in the liquid account as protection against income shocks during the second job.

2.3.2 Numerical simulations

Table 1 summarizes workers' cashout decisions at $t = 5$ in the many-period case. For each present-bias value $\beta \in \{0.4, 0.6, 0.8, 1.0\}$ and each $\sigma \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$,

[15] All other parameters remain the same as in the three-period case. I fix Social Security income y_R at 0.5.

Table 1: Cashout determinants in the many-period case

	Present-bias parameter			
	$\beta = 0.4$	$\beta = 0.6$	$\beta = 0.8$	$\beta = 1$
Intercept	2.43	1.55	1.32	1.02
	(0.02)	(0.02)	(0.01)	(0.01)
Std. dev. of income process (σ)	-0.54	1.01	0.76	0.55
	(0.01)	(0.01)	(0.01)	(0.01)
Realized avg. income before job change	-1.58	-1.51	-1.32	-1.04
	(0.02)	(0.02)	(0.01)	(0.01)

Note: This table summarizes cashout decisions made by workers in the expanded case of the model with $T = 15$. See Sections 2.1-2.3 for a description of the model. Income in each period is independently and lognormally distributed with mean 0 and standard deviation σ. For each $\sigma \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5\}$ and each present-bias parameter $\beta \in \{0.4, 0.6, 0.8, 1\}$, I conduct 10,000 simulations. The table shows regression results where the observation is a simulation and the dependent variable is an indicator equal to 1 if the worker takes a penalized 401(k) withdrawal when switching jobs at $t = 5$. The independent variables are the standard deviation of the income process σ and the worker's realized average income before the job change (i.e., $\frac{1}{5}\sum_{t=1}^{5} y_t$). I estimate the regression separately for each of the β values. Standard errors are in parentheses.

I conduct 10,000 simulations. With individual simulations as observations, I estimate regressions (separately for each β value) where the dependent variable is an indicator equal to 1 if the worker takes a penalized 401(k) withdrawal at $t = 5$. The independent variables are σ (representing the volatility of the underlying income process) and realized average income through $t = 5$ (representing the exogenous realizations of the income process).

The results in Table 1 are intuitive. First, the regression intercepts decline monotonically with β, indicating that more patient workers have lower baseline cashout rates regardless of income volatility. Second, the coefficients on the average income realization are negative and decline in magnitude with β. This is a generalization of the simpler three-period logic: workers with lower average income during their first job have less liquidity at $t = 5$ and are more likely to cash out. The relationship attenuates for higher β values because patient workers require larger negative shocks to induce cashouts (analogous to Figure 1 in the three-period case) and build larger precautionary buffers in the liquid account that provide better insurance against negative shocks (analogous to Figure 2).

The many-period case is conceptually valuable because it allows us to estimate the remaining regression coefficient on σ. This coefficient is positive for $\beta \geq 0.6$,[16] indicating that cashouts are more likely when the underlying data-generating process is more volatile (*conditional* on the average income realization through $t = 5$). This relationship holds for two reasons. The first is backward-looking: fixing realized average income, workers facing

[16]The $\beta = 0.4$ case is outlying because the worker is sufficiently impatient to cash out in *every* simulation with $\sigma = 0$. Moving to nonzero σ values introduces the possibility of large positive income shocks at $t = 5$ that deter the worker from cashing out. As a result, the regression coefficient on σ is negative.

higher σ values experience greater mean-zero volatility through $t = 5$ and are therefore more likely to deplete their liquid wealth by intertemporally smoothing consumption during the first job. The second reason is forward-looking: conditional on realized income through $t = 5$, workers facing higher σ values anticipate greater volatility during $6 \leq t \leq 10$ and thus perceive a higher continuation value for liquid wealth. The coefficient magnitudes decline with β because patient workers build larger liquidity buffers that are more likely to withstand volatility during the first job and remain sufficient for the second job.

The ceteris paribus relationships in Table 1 are instructive for the empirical analysis. The essential intuition from the three-period case persists, with negative income shocks reducing liquidity and raising cashout rates. Additionally, even if negative shocks are accompanied by positive shocks in other periods and the first-job income draws are symmetrically distributed around their expected value, the resulting mean-zero volatility can itself cause cashouts. For example, in the baseline $\beta = 1$ case with $\sigma = 0.3$, Table 1 implies that the cashout rate is 10% even in simulations where realized average income through $t = 5$ is exactly equal to its expectation.[17] The regression specifications in Section 4 will aim to capture both effects by controlling separately for the second moment (representing mean-zero volatility) and the third moment (representing asymmetric negative shocks) of participants' monthly income histories.

2.4 Key conclusions and empirical predictions

The model demonstrates that early 401(k) withdrawals can arise from both *impatience* and *illiquidity*. The impatience channel results from a standard assumption of β-δ time preferences and reproduces the well-established finding that present-biased workers make suboptimal early withdrawals from retirement accounts. The illiquidity channel (or more succinctly, the *liquidity* channel) arises because labor income is stochastic. When the worker receives negative income shocks, her current liquid resources can become small relative to her illiquid retirement wealth, making her willing to tap 401(k) accounts (and pay early-withdrawal penalties) in order to support current consumption. Liquidity-driven cashouts can occur for non-present-biased workers; in the model simulations above, even workers with $\beta = 1$ take penalized early withdrawals when they receive sufficiently large negative income shocks.

An important nuance is that the impatience and liquidity channels are not entirely

[17] With $\mathbb{E}[y_t] = e^{(1/2)*(0.3^2)}$, the cashout rate is $1.02 + (0.55 * 0.3) - (1.04 * \mathbb{E}[y_t]) = 0.097$.

distinct. Consider a present-biased worker who chooses to cash out. If a non-present-biased worker facing the same series of stochastic income draws would not have cashed out, we can say that the withdrawal was caused by impatience. But if a more favorable series of income draws would have prevented the present-biased worker from cashing out, we can also say that the withdrawal was caused by a liquidity shock. Without accurate data on 401(k) participants' time preferences and strong structural modeling assumptions, it is not possible to make fine distinctions about the proximate cause of any given cashout. The goal of the empirical analysis is instead to test the threshold hypothesis that the liquidity channel exists. If we observe an empirical relationship between income volatility and cashout probability, we can conclude that liquidity is at least a partial explanation for early withdrawals, and discuss the resulting policy implications and tradeoffs.

In addition to the key relationship between negative income shocks and early 401(k) withdrawals, the model makes two other predictions. First, because it degrades existing liquidity buffers and signals greater future liquidity need, mean-zero income volatility should be independently predictive of cashout probability. Second, the effect of income volatility should be larger for workers with lower average earnings, whose higher Social Security replacement rates put them closer to the cashout margin.[18] Since the model also demonstrates that cashout decisions are more sensitive to income volatility for lower β values, the limited evidence suggesting present bias is larger among lower-income households (Lawrance, 1991; Paserman, 2008; Fang and Silverman, 2009; Meier and Sprenger, 2015) would provide another reason to expect larger effects at the bottom of the income distribution. By assembling a high-frequency, paycheck-level income panel for 401(k) participants, I can evaluate these empirical predictions.

3 Data and summary statistics

3.1 Vanguard administrative data

I use administrative 401(k) recordkeeping data from Vanguard. The data include all information that Vanguard collects in the normal course of its recordkeeping responsibilities. At the plan level, I observe characteristics of the plan sponsor and plan rules. At the

[18] Social Security replacement rates decline steeply with lifetime average earnings, falling (in recent birth cohorts) from about 80% for very low earners to about 30% for high earners. See, e.g., Table D in Burkhalter and Chaplain (2023).

participant level, I observe all financial transactions (both contributions and withdrawals), account balances, hire and separation dates, and age.

Almost all plans in the recordkeeping data are 401(k) plans; the small remainder are highly similar but technically distinct defined-contribution plans, such as 403(b) plans. For brevity, I refer throughout the paper to the sample as being composed of "401(k) plans." I focus on a subset of plans in which I can accurately infer participants' paychecks, using the methodology described below.

3.2 Income inference

Employees contribute to 401(k) plans by deferring money from their paychecks. Usually, these payroll deferrals are specified in proportional terms: at any given time, a 401(k) participant has an "elected contribution rate" C such that C% of their next paycheck will be contributed to their 401(k) account.[19] If we can observe both the dollar amount of a given employee contribution transaction and the underlying contribution rate, we can infer the size of the corresponding paycheck. For example, if a participant with a 5% contribution rate made a $50 contribution on January 1, we can infer that they received a $1,000 paycheck on January 1. Because participants specify contribution rates as a percentage of *gross* pay, the $1,000 figure is a gross amount representing pay before tax withholding and other deductions.[20]

Table 2 illustrates the simple logic of the income inference process (using fictional transactions that do not represent any particular participant in the data). The left-hand panel shows a participant whose contribution amounts vary substantially from paycheck to paycheck while the contribution rate remains constant; the resulting variance in inferred income values likely indicates that this participant is an hourly worker. The right-hand panel shows a worker whose inferred income amounts are usually the same down to the cent, except for one large outlying transaction; this participant is likely a salaried worker who receives infrequent bonuses.

[19]The contribution rate can be set by active participant choice or by default in automatic-enrollment plans. I exclude the very small number of plans whose participants specify contributions in flat dollar terms (e.g., $50 is deferred from each paycheck, no matter the paycheck amount).

[20]This logic holds for both traditional tax-deferred and Roth contributions. In the above example, a 5% contribution rate implies $1,000 in gross pay regardless of whether the contribution is traditional or Roth. A participant making Roth contributions would have $1,000 in taxable income (rather than $950 for a participant making traditional contributions) and would therefore have higher tax withholding and lower net pay.

Table 2: Income inference examples

Likely hourly worker

Date	Contribution amount	Contribution rate	Inferred income
2022-01-07	56.54	0.06	942.33
2022-01-14	73.42	0.06	1,223.67
2022-01-21	47.60	0.06	793.33
2022-01-28	57.91	0.06	965.17
2022-02-04	60.40	0.06	1,006.67

Likely salaried worker with bonus

Date	Contribution amount	Contribution rate	Inferred income
2022-10-28	153.42	0.06	2,557.00
2022-11-11	153.42	0.06	2,557.00
2022-11-25	885.42	0.06	14,757.00
2022-12-09	178.99	0.07	2,557.00
2022-12-23	178.99	0.07	2,557.00

Note: These fictional examples are for illustrative purposes only; they do not represent any particular participant in the data.

There are two important limitations to consider when inferring income from contribution activity. First, employees' paychecks are not observable unless they contribute to their 401(k) accounts; the empirical sample is therefore limited to 401(k) *participants* and does not include non-participating employees. Even if an employee is participating in the plan, their contribution transactions may cease abruptly if they separate from employment, voluntarily set their contribution rate to zero, or hit the IRS limit on annual employee contribution amounts.[21] As discussed in more detail below, I generally restrict attention to participant-years in which contributions occur in both January and December, which helps ensure that the resulting sequence of paychecks represents the participant's complete income history for the corresponding calendar year. The main consequence of limiting to full calendar years is that I drop some high earners who hit the IRS contribution limit during the middle of the year. Since they rarely take early withdrawals, these high earners are relatively uninteresting for the main empirical analysis. I control for annual income nonparametrically throughout the empirical analysis, so dropping these high earners does not affect the estimates I obtain for low- and middle-income participants.[22]

The other limitation is that inferred paychecks do not contain information about hours worked. As a result, I cannot directly observe the hourly versus salaried nature of the participant's employment contract and cannot decompose earnings changes into wage and hours components. I address this limitation in the following ways. First, I note that the wage-hours decomposition is not relevant given the theoretical framework introduced in

[21] For example, the IRS limit on employee contributions in 2022 was $20,500. If a participant reached $20,500 in total contributions in 2022, their contributions were automatically suspended and did not restart until January 2023. Across all Vanguard-administered plans in 2022, 14% of participants reached the contribution limit. See Clark (2024).

[22] Among Vanguard-administered plans, less than 3% of participants with annual income below $100,000 hit the IRS contribution limit. See Clark (2024).

Section 2; a given earnings shock has the same effect on liquidity and cashout propensity regardless of whether it arises from a wage change or an hours change.[23] Second, I recognize that it is valuable to the empirical analysis to have at least an imperfect method for classifying participants as either hourly or salaried.

To make the binary hourly-salaried classification, I rely on statistics from the payroll-processing data used by Ganong et al. (2024b), in which workers are explicitly identified as hourly or salaried. The optimal binary-classification rule in the Ganong et al. (2024b) data is to classify a worker as hourly if at least 45% of their paychecks change by at least $1 relative to the lagged paycheck value; this rule accurately classifies about 84% of workers in their data.[24] I adopt this rule throughout the empirical analysis, classifying a participant-year as hourly if at least 45% of the participant's paychecks in the corresponding year change by at least $1. Similarly, I classify a participant as hourly if at least 45% of all of the participant's observable paychecks change by at least $1. This classification is not essential for any of the empirical analysis; I use it only as a reasonable way to split the sample and allow for heterogeneous effects of income volatility. In many cases I will refer to participants as being "likely hourly" or "likely salaried" to emphasize the imperfect nature of the classification. In cases where I simply say "hourly" and "salaried," I do so only for brevity; the classification is always based on the same 45% cutoff.

Because the main theoretical predictions from Section 2 concern the effects of negative income shocks and mean-zero volatility, infrequent positive bonus shocks among salaried workers are not of independent interest. For completeness, however, I also attempt to distinguish base from bonus income for participants who are classified as likely salaried. See Appendix B for a description of the bonus-identification algorithm. Again, to the extent that I use the bonus amounts identified by this algorithm, it is simply to divide the broader sample of salaried participants into those who do and do not receive bonuses.

3.3 Sample filtering

I study the 2008-2022 period, which is the longest time period for which all necessary components of the administrative data are available. As part of its recordkeeping responsibilities, Vanguard maintains information about the dollar amount of every contribution

[23] Additionally, Ganong et al. (2024b) show that the standard deviation of monthly earnings changes is 10 times as large as the standard deviation of monthly wage changes. It therefore seems safe to interpret earnings fluctuations within a given participant-year as arising primarily from volatility in hours worked.

[24] I thank Peter Ganong, Pascal Noel, and coauthors for generously sharing these statistics with me.

Table 3: Sample filtering steps

Filtering step	Plans	Participant-years	
		Number	Share of original
Contribution rate and W-2 income observable	1,291	20,835,949	100.0%
Contributions in Jan., Dec., and ≥ 9 distinct months	1,268	17,017,357	81.7%
Inferred annual income within 3% of W-2 value	1,224	9,382,347	45.0%

Note: The sample period is 2008-2022. I restrict to plans with observable contribution rates so that I can infer paychecks according to the methodology discussed in Section 3.2. I further restrict to plans with W-2 income data so that I can verify the sum of inferred paychecks against the annual W-2 figure.

transaction for every plan. As discussed above, I must also observe the underlying contribution rate in order to infer paychecks. Vanguard tracks participants' contribution rates for about two-thirds of the plans it administers. During the 2008-2022 period, Vanguard tracked contribution rates for 1,656 distinct plans.

For a distinct but highly overlapping subset of about two-thirds of plans, Vanguard performs annual nondiscrimination testing mandated under ERISA. As part of the testing process, these plans provide Vanguard with official annual income figures for all participants (those reported on IRS from W-2). Although the empirical analysis focuses on *sub*-annual income volatility, the annual W-2 figures are useful as a final check on the income inference process: for each participant-year, I can sum the inferred paychecks and verify that the inferred annual income amount is sufficiently close to the official W-2 figure. I therefore restrict to the 1,291 distinct plans for which Vanguard tracked contribution rates and collected W-2 income figures during the 2008-2022 period.

Table 3 shows the steps I take to arrive at the final sample of participant-years. The starting point is the universe of 20.8 million participant-years for the 1,291 eligible plans across the 2008-2022 period. After requiring that employee contribution transactions occur in January, December, and at least nine distinct months (recall that this drops some high-earning participants who hit the IRS annual contribution limit), 17 million participant-years remain.[25] Finally, I restrict to participant-years where the sum of inferred paychecks is within 3% of the official W-2 value. The final sample of 9.3 million participant-years are cases where I can be highly confident that the sequence of inferred paychecks is a complete and accurate representation of the participant's income history.

It is worth discussing the participant-years where contributions occur in January, December, and at least nine distinct months but where inferred annual income is not within

[25]A small number of participants are not paid during the summer, so I use a 9-month cutoff instead of 12-month cutoff for the purpose of identifying full calendar years.

3% of the W-2 value. Participant-years are dropped at this final filtering step for two main reasons. First, the January-December condition is a rough filter that does not perfectly identify full calendar years: for example, a participant with two paychecks in December may hit the IRS contribution limit on the first paycheck, and missing the second paycheck causes me to understate annual income by at least 3%. Second and most importantly, some participants receive non-deferrable income (i.e., income from which they are not permitted to make 401(k) contributions). Non-deferrable income is included in the official W-2 figure but by definition cannot be inferred from contribution activity. Of the 17 million participant-years with transactions in January and December, 3.2 million (19%) have non-deferrable income that is at least 3% of the W-2 value. Some plan sponsors who pay non-deferrable income do not separately report it to Vanguard, so the 3.2 million figure is a lower bound on the number of participant-years dropped for this reason.

3.4 Summary statistics

Table 4 shows summary statistics for the final sample of roughly 9.3 million participant-years. Using the binary-classification rule discussed in Section 3.2, 67% of participant-years are classified as likely hourly. Because the restriction to full calendar years drops high earners who hit the IRS annual contribution limit, this likely overstates the share of the broader recordkeeping population that is hourly. Knowing that 14% of Vanguard participants hit the contribution limit (Clark, 2024) and assuming for simplicity that all such participants are salaried, the adjusted hourly share would be 57.7%. This is roughly in line with the payroll data in Ganong et al. (2024b), where 60% of workers are hourly.[26]

The classification rule is based on a paycheck-change cutoff of 45%, but most participant-years classified as hourly have a substantially higher share of paychecks change by at least $1: the 10th percentile among the hourly group is 73%. The average hourly participant sees 92% of their paychecks in a given year change by at least $1; the corresponding average for the salaried group is 16%. Unsurprisingly, annual income is significantly lower for hourly participants: the median is $64,000 for the hourly group and $103,000 for the salaried group (in 2022 dollars). Contribution rates are more similar, averaging 7.1% for hourly participants and 7.9% for salaried participants.

To measure sub-annual income volatility, I follow Ganong et al. (2024b) in focusing on

[26]Since non-deferrable income is often awarded as bonus payments, dropping participant-years with large amounts of non-deferrable income may also disproportionately drop salaried workers and inflate the hourly share.

Table 4: Summary statistics by participant-year

	Likely hourly	Likely salary
Participant-years	6,294,909	3,087,438
Annual income (2022 $)		
10th	31,833	60,701
50th	64,106	103,222
90th	120,112	176,767
Mean	71,608	113,100
Share of paychecks with $\geq$ $1 change		
10th	73.0%	4.0%
50th	96.6%	13.1%
90th	100.0%	33.8%
Mean	91.8%	16.2%
Std. dev. of log monthly earnings/paycheck		
10th	0.05	0.01
50th	0.12	0.09
90th	0.27	0.29
Mean	0.15	0.13
Skewness of log monthly earnings/paycheck		
10th	-1.10	-1.56
50th	0.25	1.89
90th	1.69	3.01
Mean	0.26	1.26
Share negative	39.5%	26.5%
Annual contribution rate		
10th	2.8%	3.2%
50th	6.0%	6.6%
90th	12.3%	14.1%
Mean	7.1%	7.9%
Share with hardship withdrawal	2.2%	0.7%

Note: Participant-years are classified as "likely hourly" if at least 45% of the participant's paychecks in the relevant year change by at least $1 relative to the lagged paycheck value. The remaining participant-years where fewer than 45% of paychecks change by at least $1 are classified as "likely salary." See Section 3.2 for a discussion of this classification rule. Monthly income figures are normalized by the number of paychecks to eliminate mechanical variation (e.g., two- versus three-paycheck months). Annual income figures are given in real 2022 terms.

monthly earnings per paycheck. For a given participant-year, I aggregate earnings within each month and compute volatility measures (standard deviation, skewness) across the resulting sequence of 12 monthly values.[27] Normalizing each monthly earnings total by the

[27] About 3% of participant-years in the final sample have fewer than 12 months with earnings, due primarily

Table 5: Summary statistics by distinct participant

	Likely hourly	Likely salary
Distinct participants	2,035,623	1,043,633
Share auto-enrolled	36.6%	32.9%
Share separated	51.3%	55.0%
With complete spell observed	16.6%	12.3%
Share with cashout (cond. separation)	42.0%	21.2%
Cashout amount (cond. nonzero)		
Mean	12,410	17,313
Median	3,781	5,210
Cashout share of balance (cond. nonzero)		
Mean	90.8%	75.9%
Median	100.0%	100.0%
Share of cashouts for entire balance	88.0%	71.4%

Note: Participants are classified as "likely hourly" if at least 45% of their observable paychecks change by at least $1 relative to the lagged paycheck value. The remaining participants where fewer than 45% of paychecks change by at least $1 are classified as "likely salary." See Section 3.2 for a discussion of this classification rule.

number of paychecks eliminates mechanical variation arising from common pay schedules (e.g., participants paid on biweekly schedules receive three paychecks in two months per calendar year). Finally, I apply a log transformation and consider the log of monthly earnings per paycheck. This allows convenient proportional interpretations, especially when volatility measures appear in regression specifications in Section 4.

Table 4 confirms that the typical hourly participant experiences greater monthly income volatility than the typical salaried participant. For the average participant-year in the hourly group, the standard deviation of log monthly earnings per paycheck is 0.15, indicating that the typical fluctuation in monthly income is roughly 15%. For the salaried group, the corresponding average figure is 0.13. The hourly and salary figures are more similar than one might expect because roughly a third of the salaried group receive annual bonuses, creating a right tail of salaried participants with large monthly standard deviations (see the discussion of the bonus-identification algorithm in Appendix B). The difference between hourly and salaried volatility becomes clearer when looking at the third moment, i.e., the skewness of log monthly earnings per paycheck. The average hourly skewness value is 0.26, while infrequent positive bonus shocks push the corresponding salaried value up to 1.26. Hourly participants are also 1.5 times more likely to have a negative skewness value

to the small number of participants who are not paid over the summer. The sample filtering discussed in Section 3.3 ensures that all participant-years in the final sample have at least 9 months with earnings.

Table 6: Cashout rates by annual income

Annual income (2022 $)	Likely hourly	Likely salary
$\leq$ 25,000	46.4%	37.8%
25,000 – 50,000	49.5%	37.4%
50,000 – 75,000	42.2%	27.5%
75,000 – 100,000	34.9%	22.5%
100,000 – 125,000	29.9%	18.9%
125,000 – 150,000	26.2%	16.5%
$>$ 150,000	22.6%	14.7%

Note: This table considers participants who separate during the sample period of 2008-2022. Participants are classified by their average annual income across the sample period. Participants are classified as "likely hourly" if at least 45% of their observable paychecks change by at least $1 relative to the lagged paycheck value. The remaining participants where fewer than 45% of paychecks change by at least $1 are classified as "likely salary." See Section 3.2 for a discussion of this classification rule. Annual income figures are in real 2022 terms.

(40% of participant-years) than salaried participants (27%).[28] Appendix Figures A.1-A.3 show the full distributions of the standard deviation and skewness measures.

Table 5 summarizes separation and cashout outcomes for the roughly 3 million distinct participants in the final sample. For both the hourly and salaried groups, about half of participants separate from employment during the 2008-2022 sample period. Separating participants can do one of three things with their 401(k) balance: they can leave the balance in the plan, roll it over to an IRA or a new employer's 401(k) plan, or take a cash distribution. They can also combine these options, for example by taking a partial cash distribution and rolling the remaining balance to an IRA. I consider a separated participant to have cashed out if they take any nonzero cash distribution at any point after their separation date.[29]

The baseline cashout rates in Table 5 already give evidence of the model's main empirical prediction: conditional on separating, 42% of hourly participants cash out, compared to 21% of salaried participants. The large majority of cashouts are for the participant's entire vested balance, so the cashout amounts are substantial in dollar terms: the median value is $3,781 for the hourly group and $5,210 for the salaried group. Of course, hourly participants have lower annual incomes, so the difference in baseline cashout rates is driven both by the level of income and the volatility of income. Table 6, which again shows cashout rates conditional on separation, disentangles these two effects. Cashout rates decline with annual income in both the hourly and salaried groups, but for each income level, hourly cashout

[28]Salaried participants can have negative skewness values if they do not receive an annual bonus but do receive an annual raise early in the year.

[29]The sample period is 2008-2022 but the transaction data are complete through 2023, so I have a full year to observe cashout transactions for the latest cohort of separators. Cash distributions and rollovers are explicitly identified as separate transaction types in the recordkeeping data.

rates are significantly higher (12-15 percentage points higher in the $25,000-$100,000 range). The next section provides more formal tests of the volatility-cashout relationship.

4 The effect of income volatility on early withdrawals

4.1 Empirical specification

In this section, I present formal regression-based evidence on the relationship between income volatility and early 401(k) withdrawals. I consider two different withdrawal outcomes; one outcome varies at the participant-year level and the other varies at the participant level.

First, at the participant-year level, I consider the probability that a given participant takes a hardship withdrawal in a given year. Hardship withdrawals, which are generally subject to the 10% early-withdrawal penalty, allow participants experiencing financial distress to withdraw 401(k) savings during employment. Because hardship withdrawals are subject to relatively strict regulation – participants must provide documentation of a specific financial hardship and may only withdraw an amount necessary to cover the documented hardship – they are rarely used.[30] As shown in Table 4, hardship withdrawals occur in only 2.2% of hourly participant-years and 0.7% of salaried participant-years. Although they are rare, hardship withdrawals are still an informative empirical outcome because they allow me to draw a contemporaneous link between income volatility experienced in a given year and penalized 401(k) withdrawals taken in the *same* year.

The second and more important outcome is cashout probability at separation. Theoretically, post-separation cashout decisions are of primary interest because they are the explicit subject of the model discussed in Section 2. Empirically, cashouts are the most important form of leakage because they are the most common form of leakage. Again referencing Table 4, 42% of separating participants in the hourly group and 21% of separating participants in the salaried group choose to take a cash distribution. Cashouts are more prevalent than hardship withdrawals in part because they are less regulated: any participant who has separated from their employer is eligible to take a cash distribution from their 401(k) balance, without providing a specific reason or documentation of financial hardship. Cashouts are only subject to the 10% early-withdrawal penalty if the participant is younger than 55 at the end of the calendar year in which the separation occurs; I therefore

[30] See this IRS summary of hardship withdrawal rules.

restrict the baseline cashout analysis to participants who are younger than 55 at the end of their separation year.[31] Since cashouts only occur after separation, the analysis will relate cashout probability to measures of income volatility that are averaged across the participant's employment spell.

Let i index participants, j index plans, and t index years. Addtionally, let h_{ijt} be an indicator equal to one if participant i in plan j takes a hardship withdrawal in year t. The regression specification I use to study hardship withdrawals is:

$$h_{ijt} = \sum_k \gamma_k * \mathbb{1}\{e_{ijt} \in b_k\} + \sum_k \beta_k * \mathbb{1}\{e_{ijt} \in b_k\} * v_{ijt} + \rho X_{ijt} + \alpha_j + \theta_t, \tag{10}$$

where e_{ijt} is annual income, b_k are buckets for annual income, v_{ijt} is a measure of sub-annual income volatility, X_{ijt} are covariates, α_j are plan fixed effects, and θ_t are year fixed effects. The coefficients of interest are β_k, which give the effect of sub-annual income volatility on hardship withdrawal probability. The estimated effect of income volatility is conditional on nonparametric controls for annual income (via the b_k buckets), quadratics of age and plan tenure (via X_{ijt}), and the plan fixed effects (which absorb time-invariant characteristics like employer industry and plan rules). Interacting the income-volatility measure v_{ijt} with the income buckets b_k allows me to estimate heterogeneous effects across the annual income distribution. For the hardship-withdrawal analysis, I use buckets b_k that are \$5,000 wide (using annual income figures converted to real 2022 dollars).

The regression specification for the cashout analysis has the same basic structure but is specified at the participant level. Let c_{ij} be an indicator equal to one if participant i in plan j takes a nonzero cash distribution after separation. Considering only participants who separate during the sample period, the specification is:

$$c_{ij} = \sum_k \gamma_k * \mathbb{1}\{\bar{e}_{ij} \in b_k\} + \sum_k \beta_k * \mathbb{1}\{\bar{e}_{ij} \in b_k\} * \bar{v}_{ij} + \rho X_{ij} + \alpha_j + \theta_t. \tag{11}$$

Annual income $\bar{e}_{ij}$ is now averaged across all observable years for participant i in plan j, as is the volatility measure $\bar{v}_{ij}$. The age and tenure values in X_{ij} are taken as of the separation year. Because the sample size for (11) is significantly smaller (participant-years are collapsed to participants and non-separating participants are dropped), I expand the

[31] IRS rules on post-separation distributions are discussed here. I do not need to apply an age condition for the hardship-withdrawal analysis because unpenalized in-service distributions for participants older than 59.5 occur under a separate, non-hardship transaction code.

income buckets b_k to be \$10,000 wide.[32]

As v_{ijt}, I use the volatility measures introduced in Section 3.4. In some specifications, I let v_{ijt} be one-dimensional and include only the standard deviation of log monthly earnings per paycheck. In other specifications, v_{ijt} is two-dimensional and includes both the standard deviation and the skewness of log monthly earnings per paycheck. The two-dimensional case allows me to distinguish the effects of mean-zero and asymmetric volatility. Conditional on skewness, an increase in standard deviation occurs through larger mean-zero shocks; conditional on standard deviation, a decrease in skewness results from an asymmetric shift toward larger negative shocks. Finally, it is worth emphasizing for the cashout regressions that the v_{ijt} values are computed within participant-year and then averaged across years to reach $\bar{v}_{ij}$. The $\bar{v}_{ij}$ values therefore represent the average level of within-year volatility, not annual income growth across years.

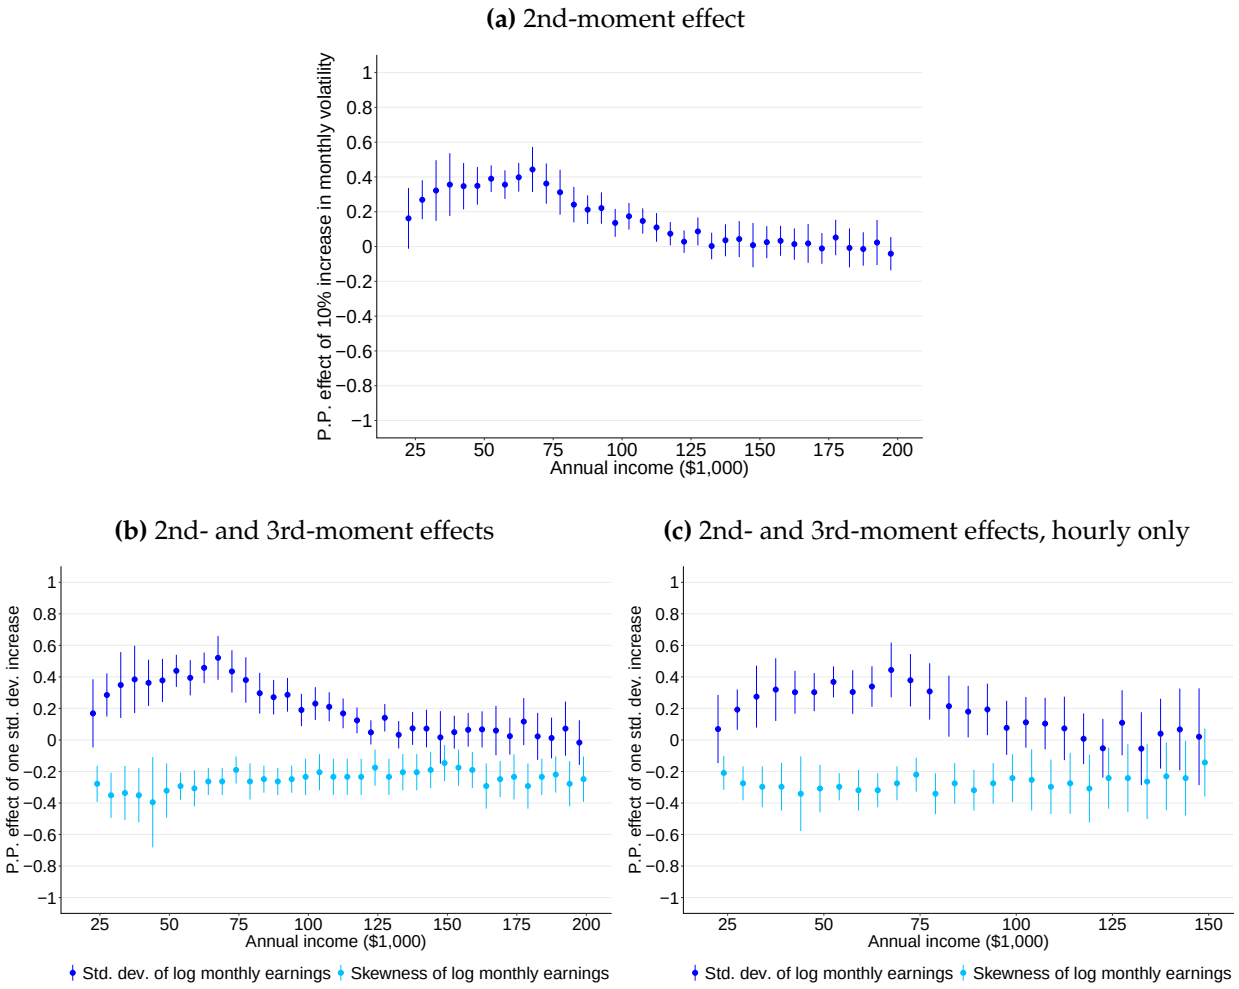
4.2 Results

Figure 3 shows the results of the hardship-withdrawal regressions from (10). The plotted coefficients are the β_k, which give the heterogeneous effects of income volatility within each annual income bucket. Panel (a) considers the entire sample of participant-years and adopts the simpler specification with only the standard deviation of log monthly earnings per paycheck included as an independent variable. I scale the coefficients so that they represent the effect of a 0.1-unit increase, which given the log transformation approximates a 10% increase in monthly income volatility. For low- and middle-income participants (up to annual income of about \$75,000), the estimated coefficients imply that a 10% increase in volatility raises hardship-withdrawal probability by about 0.4 percentage points. Given the low baseline hardship-withdrawal rate (1.7% for the overall sample), this represents a substantial proportional increase of 23%. The coefficients then decline gradually and are statistically zero for all annual income levels above \$125,000.

Panel (b) adopts the fuller specification with the skewness of log monthly earnings per paycheck included on the right-hand side. To put the standard deviation and skewness

[32] Many plans automatically execute cashout transactions for balances smaller than \$1,000 at separation and automatically execute rollover transactions for balances that are between \$1,000 and \$5,000 at separation. In (11), I therefore include dummy variables controlling for balances less than \$1,000 and for balances between \$1,000 and \$5,000 at separation. The proper econometric treatment of automatic distributions is not obvious because balance is an endogenous variable. I include the dummy variables primarily because it is conservative: balance size correlates negatively with income volatility and negatively with cashout probability. In any event, the cashout regression results are not sensitive to the inclusion or exclusion of these dummy variables.

Figure 3: The effect of income volatility on hardship-withdrawal probability

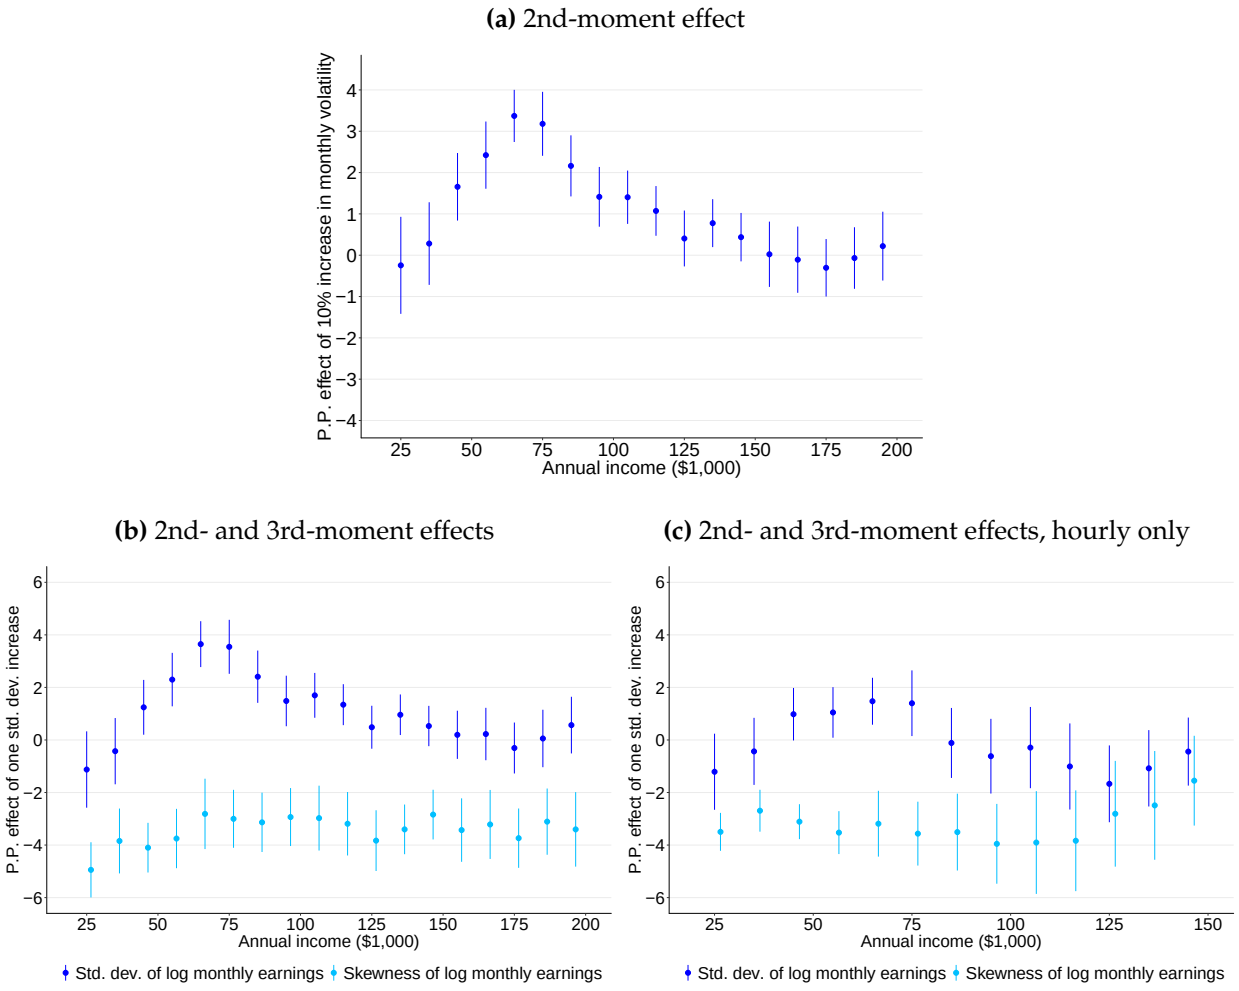


Note: Observations are participant-years and the income-volatility measures are as defined in Section 3.4. Participant-years are classified as hourly if at least 45% of the participant's paychecks in the relevant year change by at least $1 relative to the lagged paycheck value. The remaining participant-years where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. The underlying regression specification in (10) controls for annual income, age, plan tenure, plan fixed effects, and year fixed effects. The coefficient estimates in panels (b) and (c) are scaled by the standard deviation of the corresponding volatility measure. Annual income buckets on the horizontal axis are $5,000 wide and based on values in real 2022 dollars. Standard errors are clustered by plan.

measures on equal footing, I scale the coefficients so that they each represent the effect of a one-standard-deviation increase in the corresponding measure.[33] The second-moment effects are comparable to those in panel (a); the new third-moment effects are of similar magnitude but are consistently negative and persist all the way to the top of the annual

[33]This requires an unfortunate composition of standard deviations. The dark-blue dots are scaled by the standard deviation of the standard deviation of log monthly earnings per paycheck; the light-blue dots are scaled by the standard deviation of the skewness of log monthly earnings per paycheck. I compute these standard deviations over all participant-years for panel (b) and over hourly participant-years for panel (c). The same logic applies to panels (b) and (c) of Figure 4.

Figure 4: The effect of income volatility on post-separation cashout probability



Note: Observations are participants and the sample is restricted to participants who separate from employment during the sample period. I further restrict to participants for whom cashouts are penalized (i.e., those who are younger than 55 at the end of the calendar year in which they separate). The income-volatility measures are as defined in Section 3.4 and are averaged across all observable years for a given participant. Participants are classified as hourly if at least 45% of their observable paychecks change by at least $1 relative to the lagged paycheck value. The remaining participants where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. The underlying regression specification in (11) controls for average annual income during the employment spell, age at separation, plan tenure at separation, plan fixed effects, and year fixed effects. The coefficient estimates in panels (b) and (c) are scaled by the standard deviation of the corresponding volatility measure. Annual income buckets on the horizontal axis are $10,000 wide and are based on income values converted to real 2022 dollars. Standard errors are clustered by plan.

income distribution. When I restrict in panel (c) to the likely hourly group – and therefore consider only intensive-margin variation in monthly volatility among hourly workers – the results are largely the same.

Figure 4 has the same structure but reports the results of the cashout regressions from (11). In the simpler specification with only second-moment effects, the estimated coefficients rise from zero, peak at annual income values of roughly $75,000, then decline

and remain at zero above $125,000. The coefficients in panel (a) are again scaled to a 0.1-unit increase, so the peak effect implies that a 10% increase in monthly income volatility raises cashout probability by about 3.5 percentage points. Given the sample-wide average cashout rate of 34.6% among separating participants, this is a proportional increase of 10%. Again, because panel (a) considers the entire sample, these estimates are largely driven by the extensive-margin shift in volatility between the salaried and hourly groups.

Moving to panels (b) and (c) of Figure 4, the coefficients show that negative income shocks have the largest effect on cashout probability. Considering the full sample in panel (b) and focusing on the peak effect at $75,000, a symmetric perturbation that raised the second moment by one standard deviation but held fixed the third moment would increase cashout probability by about 4 percentage points. An asymmetric perturbation that raised the second moment by one standard deviation *and* lowered the third moment by one standard deviation would instead raise cashout probability by about 7 percentage points. When restricting to hourly participants (and therefore conditioning on the higher level of mean-zero volatility among this group), the second-moment effects shrink in magnitude and are at most marginally significant. Among hourly participants, the intensive-margin relationship between income volatility and cashout probability is thus primarily a third-moment effect driven by asymmetric negative shocks.

Taken together, the estimates in Figures 3 and 4 are consistent with the key model predictions from Section 2. For both in-service hardship withdrawals and post-separation cashouts, participants with more volatile income are more likely to take penalized early distributions from their 401(k) accounts. When I estimate regressions on the full sample, the extensive-margin volatility difference between the hourly and salaried groups manifests as a strong second-moment relationship between the standard deviation of monthly income and withdrawal probability. The second-moment effects validate the theoretical prediction that mean-zero volatility should independently predict cashout propensity. When I allow for both second- and third-moment effects, negative shocks that decrease the skewness of monthly income have the strongest effect on withdrawal probability. Finally, the effect of income volatility on withdrawal probability is strongest for low- and middle-income participants.

I include additional empirical results in Appendix A. Figure A.4 has the same structure as Figure 3 but the dependent variable is an indicator for 401(k) loan origination; the results are qualitatively similar and demonstrate that participants experiencing income volatility also use plan loans as a source of liquidity. Figure A.5 shows that the cashout results

hold for the subset of participants whose entire employment spells are observable in the final sample; Figure A.6 shows the results hold when including separators aged 55 and older. Finally, Figure A.7 uses as the dependent variable the share of the participant's final balance taken as a cash distribution. Because most cashouts are for the entire balance (see Table 5), the coefficient estimates are nearly identical.

5 Evidence on mechanisms

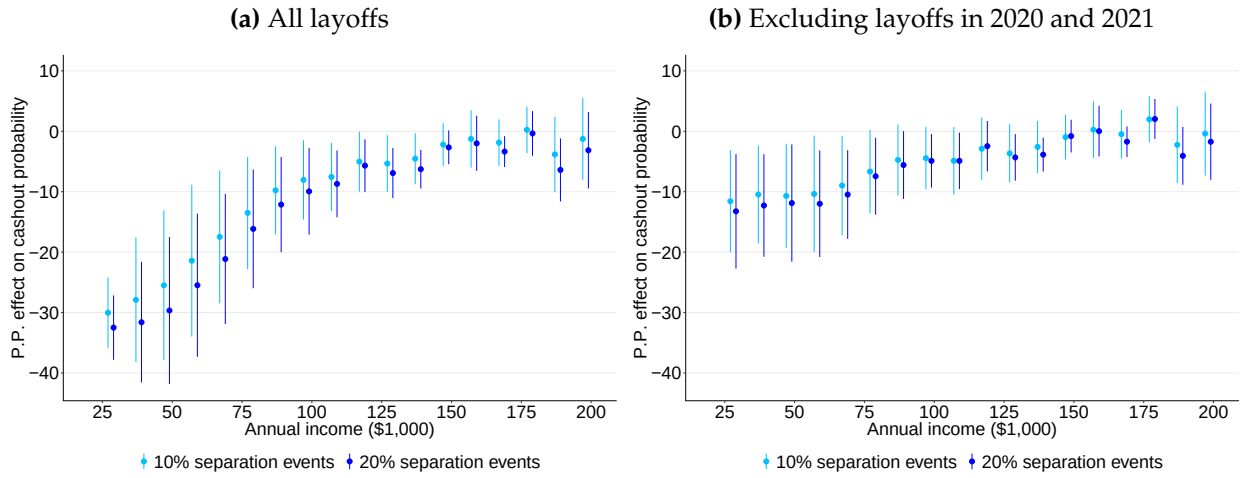
The results in the previous section establish the hypothesized relationship between income volatility and withdrawal probability. This relationship holds conditional on controls for annual income, age, tenure, and plan fixed effects. The most straightforward way to interpret the regression estimates is through the model's causal mechanism: income shocks reduce participants' current liquidity and cause them to withdraw from partially illiquid 401(k) accounts. However, recognizing that income volatility is not randomly assigned in this setting, it is worth considering whether other mechanisms could explain the results. In this section, I present two additional pieces of evidence that support the causal interpretation of income shocks acting through the liquidity channel.

5.1 Non-income liquidity: evidence from layoffs

In the model, labor income shocks cause cashouts because they reduce participants' current liquidity. If this causal liquidity mechanism underlies the empirical results, then we should observe relationships between cashout probability and other, non-labor-income factors that affect liquidity. Ideally, I would relate participants' cashout decisions to the balances in their liquid savings accounts (a^L in the model), but liquid accounts are not observable in 401(k) recordkeeping data.

Another potential source of liquidity variation comes from layoffs. When participants' separations result from involuntary layoffs, they generally become eligible for unemployment insurance (UI) benefits. If UI payments support consumption and cover essential expenses, laid-off participants may have less of an incentive to cash out their 401(k) accounts. Prior literature (Ganong and Noel, 2019) shows that household spending during unemployment is highly sensitive to UI benefit receipt, making it plausible that UI liquidity could substitute for 401(k) withdrawals. In fact, in many states, distributions from retirement accounts directly reduce the UI benefit amount for which a laid-off worker is

Figure 5: The effect of layoffs on post-separation cashout probability



Note: Observations are participants and the sample is restricted to participants who separate from employment during the sample period. I further restrict to participants for whom cashouts are penalized (i.e., those who are younger than 55 at the end of the calendar year in which they separate). I begin with the regression specification in (11) and add an indicator (interacted with the income buckets) that is equal to one if the separation is classified as a layoff. I classify the separation of participant i in plan j as a layoff if i) it is explicitly labeled as a layoff in the administrative data, or ii) at least L% of currently employed participants in plan j separate in the same month as participant i. I estimate the regressions twice, once with $L = 10\%$ and once with $L = 20\%$. In panel (b), I include separate indicators for layoffs occurring during the 2020-2021 period and those occurring outside the 2020-2021 period; I plot the latter set of non-pandemic coefficients in panel (b). The underlying regression specification in (11) controls for average annual income during the employment spell, realized income volatility during the employment spell, age at separation, plan tenure at separation, plan fixed effects, and year fixed effects. Annual income buckets on the horizontal axis are $10,000 wide and are based on income values converted to real 2022 dollars. Standard errors are clustered by plan.

eligible, making 401(k) withdrawals an ineffective source of post-layoff liquidity.[34] Given that UI replacement rates are higher for lower-income workers and were significantly increased by federal supplements during the pandemic period (Ganong et al., 2020, 2024a), we can make two supplemental predictions: the negative effect of layoffs on cashout probability should be larger at the bottom of the income distribution and during the 2020-2021 period.

The administrative recordkeeping data contain a field giving the reason for participant separations, but some plan sponsors do not report separation reasons to Vanguard. I therefore use a combination of direct information and indirect inference to identify layoffs. I consider the separation of participant i in plan j to be a layoff if i) the separation is explicitly identified as a layoff in the administrative data, or ii) at least L% of currently employed participants in plan j separate in the same month as participant i, for some chosen threshold L. The logic behind the second condition is that tightly bunched separations

[34]For example, in Pennsylvania, as long the 401(k) balance contains a nonzero amount of employer contributions, weekly UI benefits are reduced by 50% of the pro-rated weekly amount of any 401(k) withdrawal taken after the layoff.

typically indicate firm-wide layoffs rather than correlated voluntary separations (Jacobson et al., 1993).[35]

Figure 5 shows the results of the layoff analysis. I begin with the cashout regression specification in (11) and add an indicator (interacted with each annual income bucket) that that is equal to one if the participant's separation is classified as a layoff. I estimate the regression with two different cutoffs for the plan-level monthly separation rate ($L = 10\%$ and $L = 20\%$) and report both sets of results. Because the specification in (11) already includes the volatility measures $\bar{v}_{ij}$ (i.e, the standard deviation and skewness of log monthly earnings per paycheck), the layoff coefficients give the marginal effect of the separation being a layoff, conditional on realized income volatility during the employment spell.

Consistent with the causal liquidity mechanism, the layoff indicators are consistently negative and are largest at the bottom of the annual income distribution. Becuase relatively few separations are classified as layoffs (10.7% of separations with $L = 10\%$ and 8.9% of separations with $L = 20\%$), the standard errors are relatively large but still tight enough to make estimates below the $100,000 annual income level statistically significant. At the lowest annual income levels, the estimates are quite large (about 30 percentage points), implying that layoffs take average low-income cashout rates down from 40% to about 10% (see Table 6). Finally, in panel (b), I include separate indicators for layoffs occurring during the 2020-2021 period and for those occurring outside the 2020-2021 period, and plot the latter set of non-pandemic coefficients (the small sample size makes the pandemic coefficients themselves highly imprecise). As expected, the negative layoff effect is smaller outside the 2020-2021 period, when UI benefits were less generous.[36]

5.2 Labor demand versus supply: IV evidence

The model assumes that the labor income process y_t is exogenous. Similarly, the regression estimates in Section 4 have a causal interpretation only if the income-volatility measures v_{ijt} are uncorrelated with participant-level unobservables that determine withdrawal probability. The exogenous-income assumption is most accurate when income volatility arises from labor demand shocks – for example, when a firm experiencing a slow period cuts shifts for all of its hourly workers in a given month. If income shocks arise

[35] I only apply the second condition to plans with at least 1,000 participants.

[36] The negative layoff coefficients here are in contrast to Goodman et al. (2021), who find in federal tax-return data that separators who receive UI payments are *more* likely to take penalized DC withdrawals than other separators.

instead from labor supply factors, then the estimates should be interpreted with more care.

Some supply-driven income shocks are still consistent with the liquidity channel of early withdrawals. For example, consider a participant who misses shifts because of illness, then goes to the doctor and incurs medical expenses. The missed shifts would appear in the data as a negative monthly income shock, and if the participant takes a subsequent 401(k) withdrawal to cover the medical expenses, the exogeneity assumption for v_{ijt} is violated. Even so, correlated income and expense shocks are consistent with the liquidity channel because the participant takes an early withdrawal to cover a consumption need that they cannot meet with their current liquid resources. The instigating factor is still a liquidity shock, but the monthly income data understate the size of the shock.

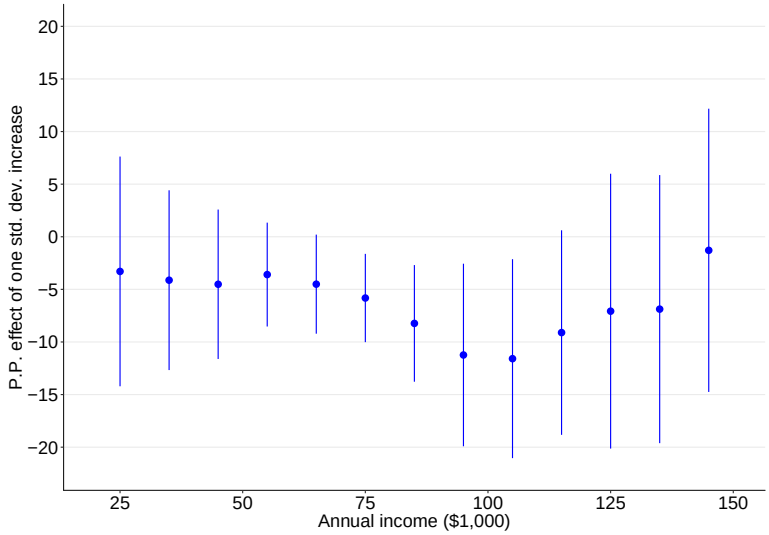
Other types of labor supply volatility pose a greater identification threat. Consider a second earner or part-time worker who may have substantial retirement assets outside the 401(k) plan (e.g., from a higher-earning spouse). This worker might be less attached to their job and may therefore selectively accept shifts, take unpaid vacations, or otherwise display a volatile monthly earnings history. At the same time, because their 401(k) balance is small relative to other sources of household wealth, they may feel less compelled to preserve it and may be more likely to cash out at separation. I address the possibility of endogenous income volatility among low-attachment workers in two ways.

First, I cite prior work showing that unpaid vacations are not an important source of worker-level income volatility. In one of the few existing studies using paycheck-level earnings data, Ganong et al. (2024b) develop an algorithm to predict unpaid vacation spells and show that unpaid vacation explains a minimal share of monthly income volatility. If unpaid vacations are similarly unimportant among the present sample of 401(k) participants, then high withdrawal probabilities among marginally attached workers becomes less of an identification concern. I also note that the effect of income volatility on early withdrawals is present for a wide range of annual income levels, not just for the lowest earners who are less likely to be strongly attached to their jobs (see again Figures 3-4).

Second, I use an instrumental-variable design to isolate monthly income volatility that arises from firm-wide labor demand shocks. Because the IV estimates (though noisier) are similar to the baseline estimates from Section 4, I conclude that endogenous labor supply is likely not driving the results.

Focusing on post-separation cashouts, I use coworkers' income volatility as a demand-driven instrument for own volatility (Koustas, 2018; Ganong et al., 2023). In particular,

Figure 6: The effect of monthly income skewness on cashout probability: IV estimates



Note: Observations are participants and the sample is restricted to participants who separate from employment during the sample period. I further restrict to participants for whom cashouts are penalized (i.e., those who are younger than 55 at the end of the calendar year in which they separate) and who are classified as likely hourly (see Section 3.2 for a discussion of the classification rule). The independent variable is the skewness of log monthly earnings per paycheck, averaged over the years comprising the participant's employment spell (see Section 3.2 for a discussion of monthly income volatility measures). I instrument for participant i's earnings skewness using the average earnings skewness of coworkers at participant i's plan during participant i's employment spell (see Section 5.2 for a discussion of the IV estimation process). I estimate the IV equations separately for each annual income bucket and control for age at separation, plan tenure at separation, and year fixed effects. The coefficient estimates are scaled by the standard deviation of monthly income skewness across the sample of likely hourly participants. Standard errors are clustered by plan.

I estimate the following set of equations:

$$\hat{v}_{ij} = \alpha_0 + \alpha_1 \bar{v}_{i'(j)\neq i} + \alpha_2 X_{ij}, \tag{12}$$

$$c_{ij} = \gamma_0 + \gamma_1 \hat{v}_{ij} + \gamma_2 X_{ij}. \tag{13}$$

Here, as in the baseline regression specification in (11), $\bar{v}_{ij}$ is a measure of monthly income volatility for separated participant i in plan j, averaged over the years comprising participant i's employment spell. I obtain the first-stage predicted value $\hat{v}_{ij}$ by regressing $\bar{v}_{ij}$ on the leave-out coworker mean, $\bar{v}_{i'(j)\neq i}$, which is similarly computed over the years comprising participant i's employment spell. The second stage regresses the cashout indicator c_{ij} on the first-stage predicted values, with γ_1 the coefficient of interest. The exogenous covariates X_{ij} include quadratics of age and plan tenure at separation and year fixed effects.

Since I only expect hourly workers' income to be responsive to monthly firm labor demand shocks, I restrict the IV sample (both the separated participants i and the coworker pool $i'(j) \neq i$) to participants who are classified as likely hourly. Because negative income shocks are of primary interest, I let the volatility measure $\bar{v}_{ij}$ be the skewness of log monthly

earnings per paycheck. The first-stage relationship is strong: regressing own on coworkers' monthly income skewness yields a slope of 0.99 with a standard error of 0.0008 and an R-squared value of 0.12 (see Figure A.8).

Figure 6 shows the IV estimates. To obtain heterogeneous treatment effects comparable to those in Section 4, I estimate (12)-(13) separately for each annual income bucket and plot the resulting set of second-stage coefficients. Because the first-stage relationship with plan-level coworker skewness is almost one-for-one, the variation predicted by the instrument is essentially at the plan-year level. As a result, the standard errors are significantly larger than for the corresponding baseline skewness estimates in Figure 3c. Even so, the point estimates are quite similar and if anything are slightly larger in magnitude for higher annual income levels. The IV estimates thus give a relatively noisy confirmation that the key cashout results hold when isolating negative income shocks to those arising from firm-wide demand factors.

6 Discussion and policy implications

The key conclusion from the preceding empirical results is that the liquidity channel exists. Of course, the liquidity channel can coexist with other explanations for early withdrawals, and the motivating model from Section 2 explicitly accounts for the possibility of present bias. Impatience likely remains an important explanation for leakage, and policymakers therefore face a tradeoff: higher early-withdrawal penalties mitigate behavioral biases but impose additional costs on liquidity-constrained households.

This tradeoff has been studied extensively in the literature on optimal retirement policy (Beshears et al., 2025; Amador et al., 2006; Beshears et al., 2015; Laibson et al., 1998). The theoretical conclusion from this literature is that a bifurcated two-account system is, or is nearly, optimal: one account should be completely liquid and the other completely illiquid (Beshears et al., 2025; Amador et al., 2006). In these models, the illiquid account (often resembling the U.S. Social Security system) prevents early withdrawals among present-biased households, while the liquid account allows a limited amount of flexibility for patient households with variable pre-retirement consumption preferences. The current 10% penalty makes 401(k) accounts neither completely liquid nor completely illiquid, and so a natural policy question is whether 10% is too high or too low.

Outside of highly stylized settings with two or three periods, optimal liquidity policy is difficult to characterize. Even in the simple three-period model considered here, precise

quantitative statements about optimal withdrawal rules would depend on empirical objects that are not currently observable alongside 401(k) recordkeeping data: the joint distribution of present bias and income volatility, the volatility of household (rather than individual) income, and access to credit. The distribution of present bias (and in particular whether this distribution is wide or narrow relative to the variance in marginal utility arising from income shocks) is a crucial model input, but the empirical distribution of time preferences is still largely an open question.[37]

Setting aside the question of whether 10% is the optimal uniform withdrawal penalty, the empirical results suggest that there could be value in targeting limited amounts of 401(k) liquidity at participants who are more likely to have true liquidity need. Two such liquidity policies recently became available to plan sponsors as optional 401(k) features permitted under the Setting Every Community Up for Retirement Enhancement (SECURE) Act of 2022. First, plan sponsors may now permit penalty-free emergency withdrawals. At adopting plans, participants may withdraw up to $1,000 per year for emergency liquidity purposes without paying the 10% early withdrawal penalty.[38] If emergency withdrawals help participants meet liquidity needs during their employment spells, they may have stronger liquidity positions at separation and may be less likely to cash out their balances when switching jobs.

The second liquidity feature newly permitted under the SECURE Act is the Pension-Linked Emergency Savings Account (PLESA). At adopting plans, participants with annual incomes below an IRS-specified cutoff can make Roth payroll deferrals into emergency savings accounts, from which they can withdraw at any time without penalty. Employee contributions to PLESAs earn matching contributions according to the employer's 401(k) match schedule, so participants need not forgo matching dollars to fund the emergency account.[39] Again, if PLESAs help low- and middle-income participants meet liquidity needs during employment, they may prevent some penalized, full-balance cashouts and reduce net leakage.[40]

[37] See the literature reviews in Lockwood (2020) and Cohen et al. (2020). The relevant empirical object is the distribution of present bias *among 401(k) participants*, whom we might expect to be more patient than the population at large. Automatic enrollment may induce some relatively impatient workers to participate in 401(k) plans, but given the 2008-2022 sample period, only about a third of the empirical sample here was automatically enrolled (see Table 5).

[38] The documentation requirements for the new emergency withdrawals are also more relaxed than for traditional hardship withdrawals, as plan sponsors can rely on participants' self-certification of financial need. See Beshears et al. (2024b) for a detailed discussion of the emergency-withdrawal feature.

[39] See this IRS summary of PLESA rules.

[40] Some employers view the rules governing PLESAs as too complex and instead offer simpler emergency

In general, whether through PLESAs or other employer-sponsored programs, it seems promising to leverage the same workplace-based tools and incentives that have boosted retirement saving (automatic enrollment, automatic payroll deferrals, matching contributions) to encourage liquid precautionary saving among low- and middle-income workers. The substantial monthly income volatility I document among 401(k) participants makes precautionary wealth valuable, and the model demonstrates that the same behavioral biases restraining retirement saving also lead workers to undersave in liquid accounts (see again Figure 2). Non-retirement savings programs are still relatively rare among U.S. firms, but recent evidence from the U.K. indicates that automatic enrollment into employer-sponsored short-term savings accounts significantly increases participation rates and liquid wealth accumulation (Berk et al., 2023).

7 Conclusion and opportunities for future work

Using paycheck-level income data for a large sample of 401(k) participants, I demonstrate a systematic relationship between workers' monthly income volatility and their propensity to take penalized pre-retirement withdrawals. The empirical results are consistent with a simple model in which workers use 401(k) assets to support consumption when income shocks reduce their current liquid resources. The liquidity channel of early withdrawals is distinct from, and likely exists alongside, the impatience channel arising from present-biased time preferences. As a result, 401(k) withdrawal rules involve a tradeoff between mitigating present bias and providing flexibility to liquidity-constrained households.

There are several promising opportunities for future work. First, linking 401(k) data with other household financial information would permit a fuller analysis of withdrawal decisions. I focus on income volatility because I can infer paychecks from recordkeeping data, but the model shows that liquid account balances are also an important determinant of cashout behavior. Data on liquid accounts would allow us to verify that participants who choose to cash out have low liquid wealth, and would permit joint analysis of participants' liquid and illiquid saving decisions. Similarly, linking 401(k) data with credit reports would allow us to determine whether income shocks lead to accumulation of costly credit card

savings vehicles outside the 401(k) plan; the potential liquidity benefits of these non-401(k) emergency accounts are similar. See Ann Carrns, "Employers Can Now Enroll Workers in Some Emergency Savings Accounts," *The New York Times*, February 9, 2024.

debt and whether some cashouts are used to pay off debt.

Second, survey data would provide a valuable complement to the observational analysis undertaken here. Direct survey evidence on the factors motivating cashouts would allow an independent test of the liquidity hypothesis alongside other potential explanations (like impatience or uncertainty about taxes and penalties).

Finally, analysis of 401(k) leakage – and retirement policy analysis more broadly – would benefit from additional evidence on the empirical distribution of time preferences. The welfare consequences of early-withdrawal penalties depend heavily on the average level of present bias, the degree of heterogeneity across 401(k) participants, and potential correlations with characteristics like age, annual income, and income volatility. As the empirical literature on time preferences grows, quantitative policy analysis can become more precise.

References

AMADOR, M., I. WERNING, AND G.-M. ANGELETOS (2006): "Commitment vs. Flexibility," *Econometrica*, 74, 365–396.

AMERIKS, J., A. CAPLIN, J. LEAHY, AND T. TYLER (2007): "Measuring Self-Control Problems," *American Economic Review*, 97, 966–972.

AMROMIN, G. AND P. SMITH (2003): "What Explains Early Withdrawals from Retirement Accounts? Evidence from a Panel of Taxpayers," *National Tax Journal*, 56, 595–612.

ARGENTO, R., V. L. BRYANT, AND J. SABELHAUS (2015): "Early Withdrawals from Retirement Accounts During the Great Recession," *Contemporary Economic Policy*, 33, 1–16.

BERK, S. H., J. J. CHOI, J. GARG, J. BESHEARS, AND D. LAIBSON (2023): "Automating Short-Term Payroll Savings: Evidence from Two Large U.K. Experiments," Working Paper.

BESHEARS, J., M. BLAKSTAD, J. J. CHOI, C. FIRTH, J. GATHERGOOD, D. LAIBSON, R. NOTLEY, J. D. SHETH, W. SANDBROOK, AND N. STEWART (2024a): "Does Pension Automatic Enrollment Increase Debt? Evidence from a Large-Scale Natural Experiment," Working Paper.

BESHEARS, J., J. J. CHOI, C. CLAYTON, C. HARRIS, D. LAIBSON, AND B. C. MADRIAN (2025): "Optimal Illiquidity," *Journal of Financial Economics*, 165.

BESHEARS, J., J. J. CHOI, J. DICKSON, A. GOODMAN, F. GREIG, AND D. LAIBSON (2024b): "Does 401(k) Loan Repayment Crowd Out Retirement Saving? Implications for Plan Design," Wharton Pension Research Council Working Paper No. 2024-22.

BESHEARS, J., J. J. CHOI, J. HURWITZ, D. LAIBSON, AND B. C. MADRIAN (2015): "Liquidity in Retirement Saving Systems: An International Comparison," *American Economic Review*, 105, 420–425.

BESHEARS, J., J. J. CHOI, D. LAIBSON, AND B. C. MADRIAN (2009): "The Importance of Default Options for Retirement Savings Outcomes: Evidence from the United States," in *Social Security Policy in a Changing Environment*, ed. by J. R. Brown, University of Chicago Press, 167–198.

BESHEARS, J., J. J. CHOI, D. LAIBSON, AND P. MAXTED (2022): "Present Bias Causes and Then Dissipates Auto-enrollment Savings Effects," *American Economic Association Papers and Proceedings*, 112, 136–141.

BURKHALTER, K. AND C. CHAPLAIN (2023): "Replacement Rates for Hypothetical Retired Workers," SSA Actuarial Note 2023.9.

BUSCH, C., D. DOMEIJ, F. GUVENEN, AND R. MADERA (2022): "Skewed Idiosyncratic Income Risk Over the Business Cycle: Sources and Insurance," *American Economic Journal: Macroeconomics*, 14, 207–242.

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.

CHOI, J. J., D. LAIBSON, AND B. C. MADRIAN (2004): "Plan Design and 401(k) Savings Outcomes," *National Tax Journal*, 57, 275–298.

CHOI, J. J., D. LAIBSON, B. C. MADRIAN, AND A. METRICK (2002): "Defined Contribution Pensions: Plan Rules, Participant Choices, and the Path of Least Resistance," in *Tax Policy and the Economy*, ed. by J. M. Poterba, MIT Press, 67–113.

CHOUKHMANE, T., J. COLMENARES, C. O'DEA, J. ROTHBAUM, AND L. D. SCHMIDT (2024): "Who Benefits from Retirement Savings Incentives in the U.S.? Evidence on Gaps in Retirement Wealth Accumulation by Race and Parental Income," Working Paper.

CHOUKHMANE, T. AND C. PALMER (2024): "How Do Consumers Finance Increased Retirement Savings?" Working Paper.

CLARK, J. W. (2024): "How America Saves 2024," Vanguard Report.

COHEN, J., K. M. ERICSON, D. LAIBSON, AND J. M. WHITE (2020): "Measuring Time Preferences," *Journal of Economic Literature*, 58, 299–347.

FANG, H. AND D. SILVERMAN (2009): "Time-Inconsistency and Welfare Program Participation," *International Economic Review*, 50, 1043–1077.

FARRELL, D. AND F. GREIG (2015): "Weathering Volatility: Big Data on the Financial Ups and Downs of U.S. Individuals," JPMorgan Chase Institute Report.

FARRELL, D., F. GREIG, AND C. YU (2019): "Weathering Volatility 2.0: A Monthly Stress Test to Guide Savings," JPMorgan Chase Institute Report.

FULFORD, S. L. AND D. LOW (2024): "Expense Shocks Matter," Working Paper.

GANONG, P., F. GREIG, P. NOEL, D. M. SULLIVAN, AND J. VAVRA (2024a): "Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data," *American Economic Review*, 114, 2898–2939.

GANONG, P., D. JONES, P. NOEL, D. FARRELL, F. GREIG, AND C. WHEAT (2023): "Wealth, Race, and Consumption Smoothing of Typical Labor Income Shocks," Working Paper.

GANONG, P. AND P. NOEL (2019): "Consumer Spending During Unemployment: Positive and Normative Implications," *American Economic Review*, 109, 2383–2424.

GANONG, P., P. NOEL, C. PATTERSON, J. VAVRA, AND A. WEINBERG (2024b): "Earnings Instability," Working Paper.

GANONG, P., P. NOEL, AND J. VAVRA (2020): "US Unemployment Insurance Rates During the Pandemic," *Journal of Public Economics*, 191.

GOODMAN, A. C., J. GOODMAN, L. GOODMAN, AND S. GOODMAN (2015): "A Few Goodmen: Surname-Sharing Economist Coauthors," *Economic Inquiry*, 53, 1392–1395.

GOODMAN, L., J. MORTENSEN, K. MACKIE, AND H. R. SCHRAMM (2021): "Leakage from Retirement Savings Accounts in the United States," *National Tax Journal*, 74, 689–719.

GUVENEN, F., F. KARAHAN, S. OZKAN, AND J. SONG (2021): "What do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" *Econometrica*, 89, 2303–2339.

GUVENEN, F., S. OZKAN, AND J. SONG (2014): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122, 621–660.

JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): "Earnings Losses of Displaced Workers," *American Economic Review*, 83, 685–709.

KOUSTAS, D. K. (2018): "Consumption Insurance and Multiple Jobs: Evidence from Rideshare Drivers," Working Paper.

LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–478.

LAIBSON, D., A. REPETTO, J. TOBACMAN, R. E. HALL, W. G. GALE, AND G. A. AKERLOF (1998): "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity*, 1, 91–196.

LAWRANCE, E. (1991): "Poverty and the Rate of Time Preference: Evidence from Panel Data," *Journal of Political Economy*, 99, 54–77.

LOCKWOOD, B. B. (2020): "Optimal Income Taxation with Present Bias," *American Economic Journal: Economic Policy*, 12, 298–327.

MADRIAN, B. C. AND D. F. SHEA (2001): "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics*, 116, 1149–1187.

MEEUWIS, M., D. PAPANIKOLAOU, J. ROTHBAUM, AND L. D. SCHMIDT (2025): "Time-Varying Risk Premia and Heterogeneous Labor Market Dynamics," Working Paper.

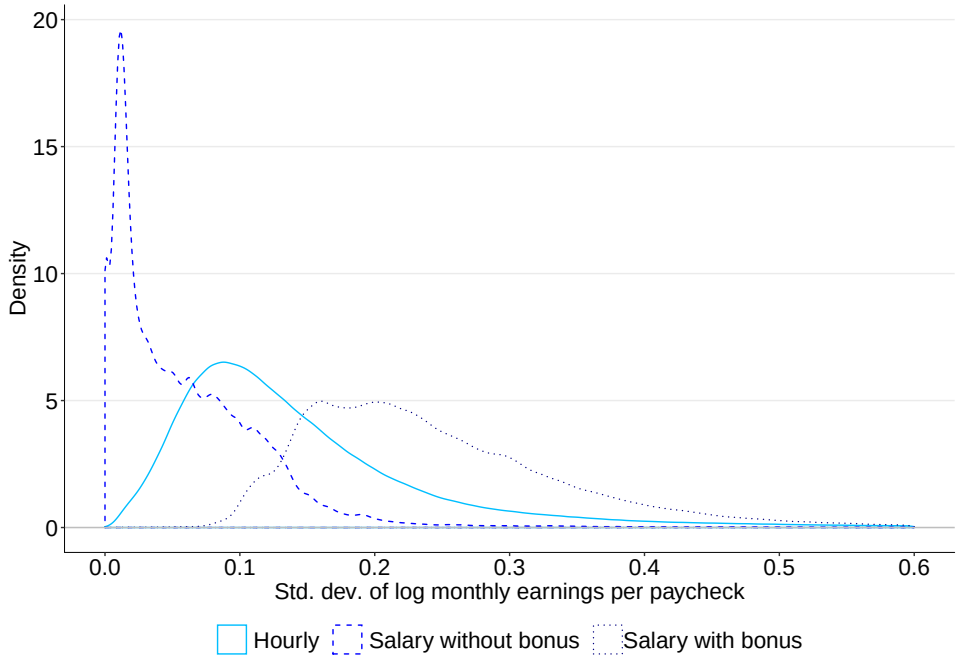
MEIER, S. AND C. D. SPRENGER (2015): "Temporal Stability of Time Preferences," *Review of Economics and Statistics*, 97, 273–286.

MOSER, C. AND P. O. DE SOUZA E SILVA (2019): "Optimal Paternalistic Savings Policies," Working Paper.

PASERMAN, M. D. (2008): "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation," *The Economic Journal*, 118, 1418–1452.

Appendix A Additional empirical results

Figure A.1: The standard deviation of monthly income, within participant-year



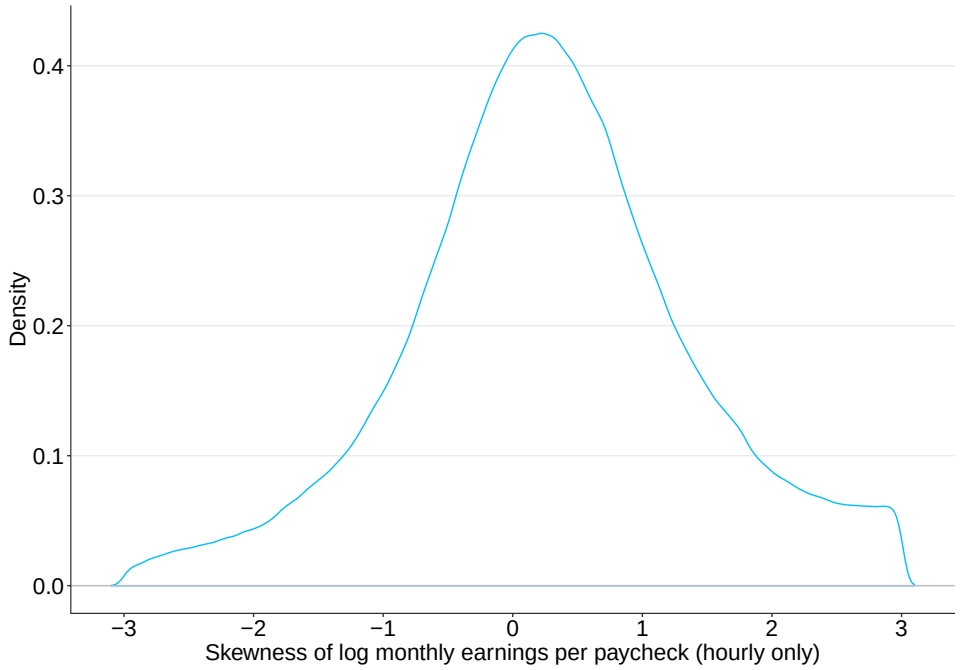
Note: Observations are participant-years; volatility measures computed over monthly values within a given participant-year. Participant-years are classified as hourly if at least 45% of the participant's paychecks in the relevant year change by at least $1 relative to the lagged paycheck value. The remaining participant-years where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. Bonuses are identified using the algorithm discussed in Appendix B.

Figure A.2: The skewness of monthly income, within participant-year



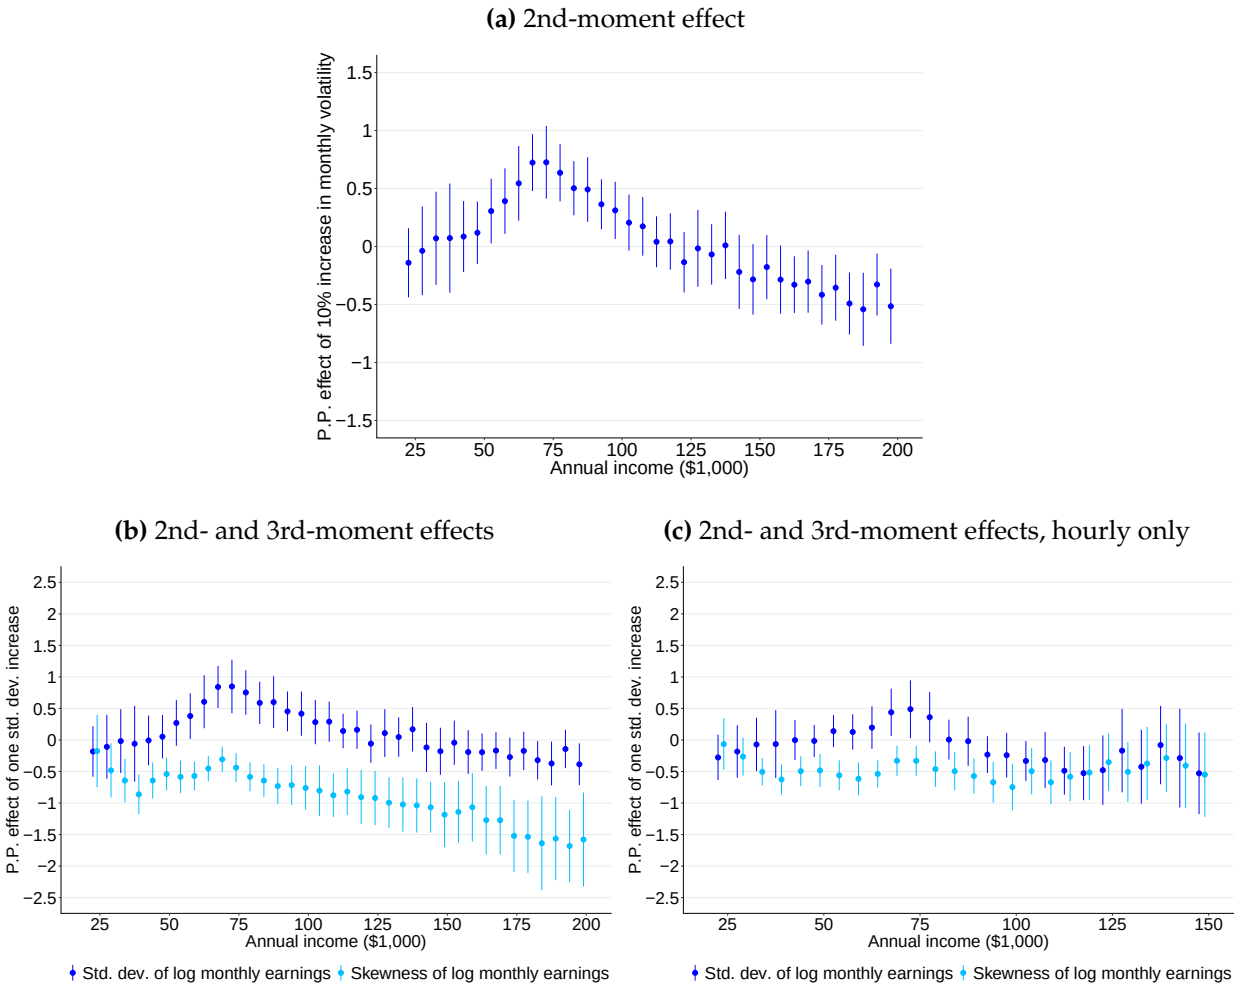
Note: Observations are participant-years; volatility measures computed over monthly values within a given participant-year. Participant-years are classified as hourly if at least 45% of the participant's paychecks in the relevant year change by at least $1 relative to the lagged paycheck value. The remaining participant-years where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. Bonuses are identified using the algorithm discussed in Appendix B. If X is a vector of 12 values, 11 of which are identical and strictly less than the 12th, then the skewness of X is approximately 3. Salaried participants with bonuses (who often have 11 very similar monthly income values and an outlying 12th) therefore bunch around 3.

Figure A.3: The skewness of monthly income, within participant-year (hourly only)



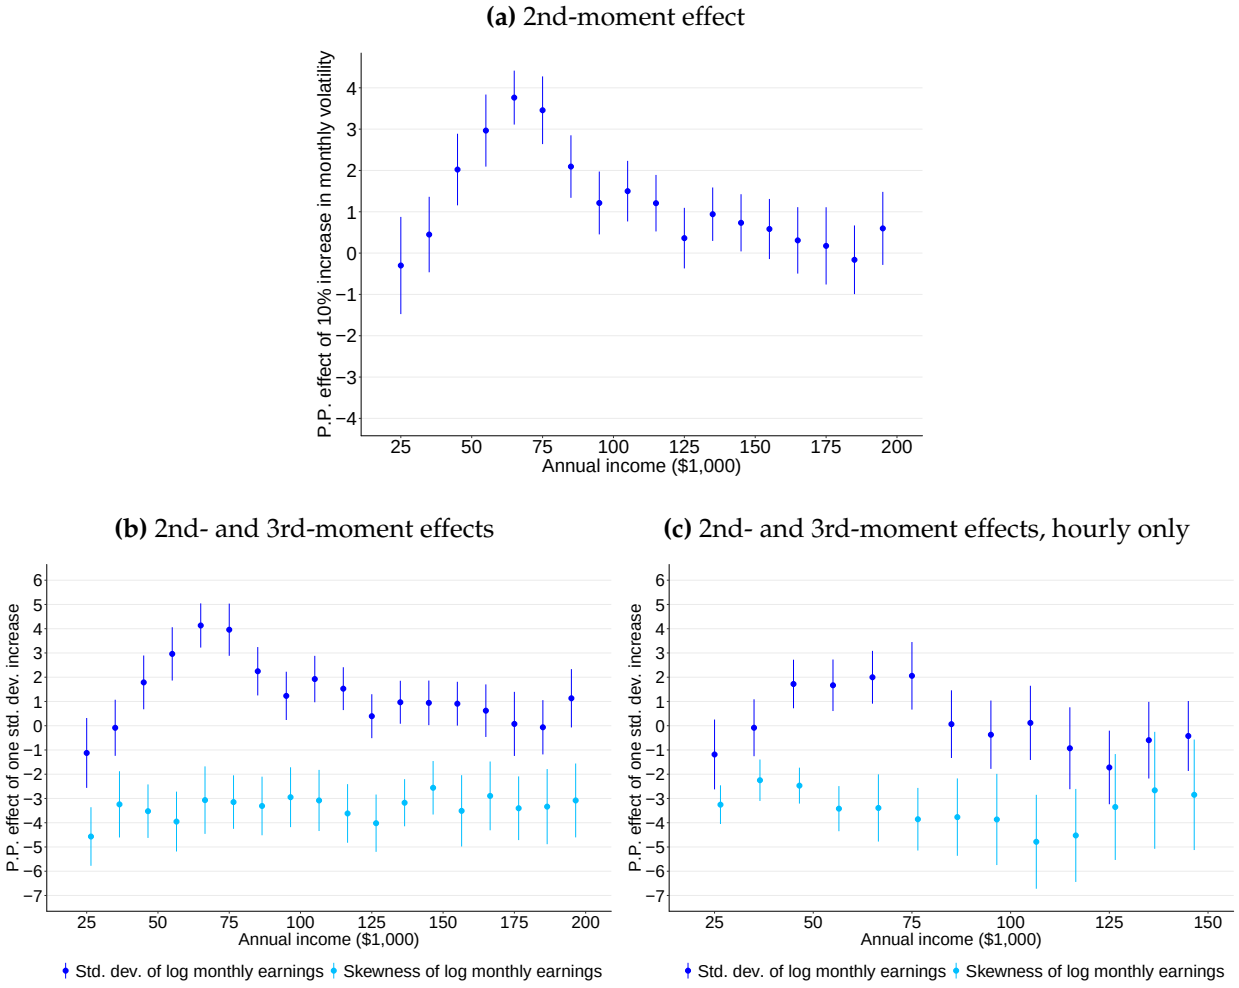
Note: Observations are participant-years; volatility measures computed over monthly values within a given participant-year. Participant-years are classified as hourly if at least 45% of the participant's paychecks in the relevant year change by at least $1 relative to the lagged paycheck value. The remaining participant-years where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule.

Figure A.4: The effect of income volatility on loan-origination probability



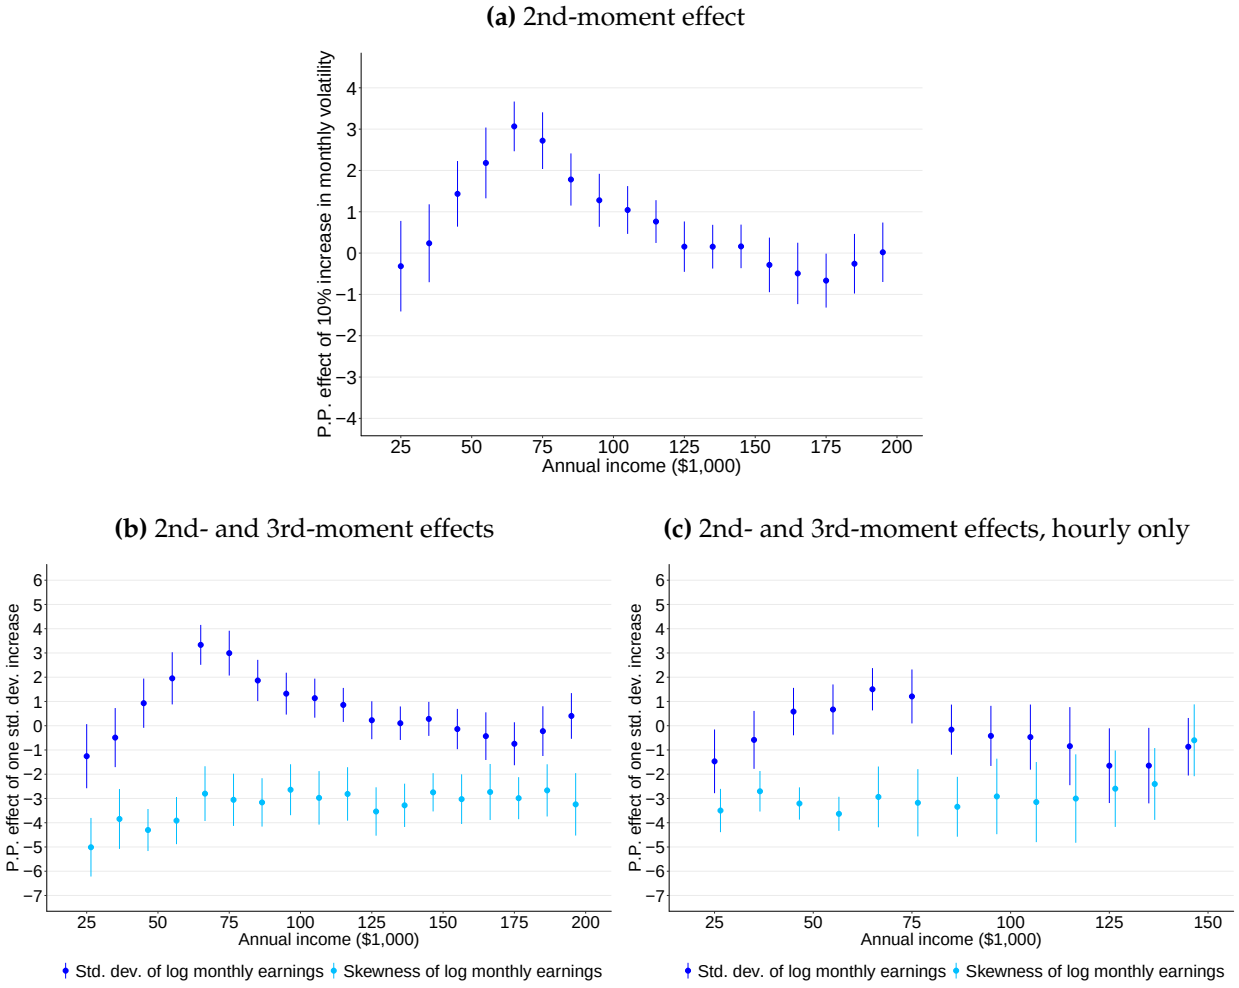
Note: Observations are participant-years and the income-volatility measures are as defined in Section 3.4. The dependent variable is an indicator equal to one if the participant originates a 401(k) plan loan in the relevant year. Plan loans allow participants to borrow against the vested balance in their 401(k) account, paying back the account by making additional payroll deferrals over a prespecified repayment period. If a participant fails to make the repayments on schedule, the remaining loan balance is considered a cash distribution on which the participant must pay income taxes and a 10% penalty (the penalty only applies to participants younger than 59.5). Participant-years are classified as hourly if at least 45% of the participant's paychecks in the relevant year change by at least $1 relative to the lagged paycheck value. The remaining participant-years where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. The underlying regression specification in (10) controls for annual income, age, plan tenure, plan fixed effects, and year fixed effects. The coefficient estimates in panels (b) and (c) are scaled by the standard deviation of the corresponding volatility measure. Annual income buckets on the horizontal axis are $5,000 wide and based on income values converted to real 2022 dollars. Standard errors are clustered by plan.

Figure A.5: The effect of income volatility on post-separation cashout probability, restricting to balanced panel of participant-years



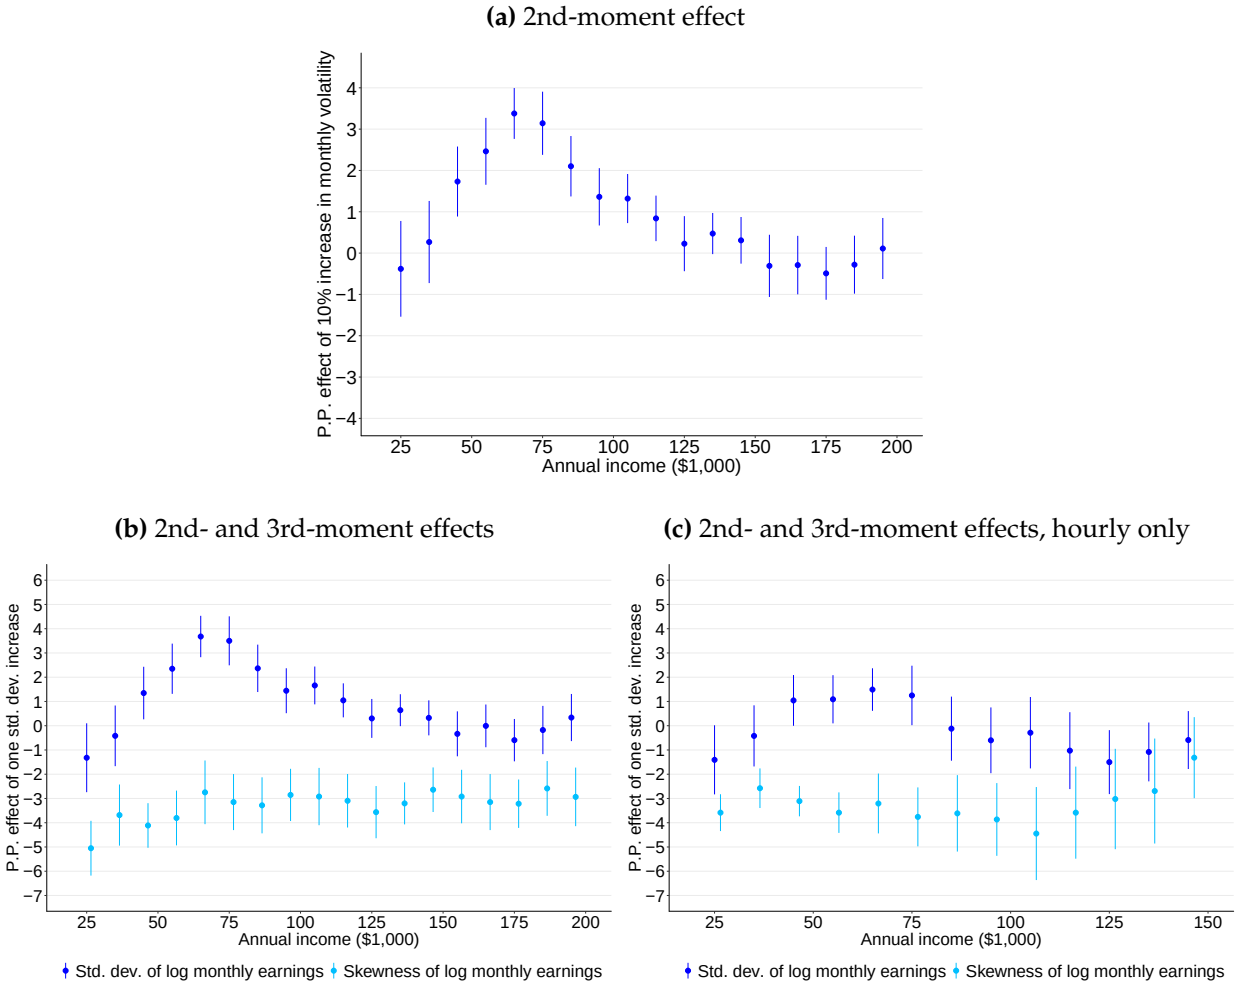
Note: Observations are participants and the sample is restricted to participants who separate from employment during the sample period. I further restrict to participants for whom cashouts are penalized (i.e., those who are younger than 55 at the end of the calendar year in which they separate). Unlike the baseline analysis in Figure 4, I also restrict to participants who appear in the final sample of participant-years for their entire employment spell (i.e., they begin and separate from employment within the sample period and their inferred annual income is within 3% of the corresponding W-2 value for each year of their employment spell). The income-volatility measures are as defined in Section 3.4 and are averaged across all observable years for a given participant. Participants are classified as hourly if at least 45% of their observable paychecks change by at least $1 relative to the lagged paycheck value. The remaining participants where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. The underlying regression specification in (11) controls for average annual income during the employment spell, age at separation, plan tenure at separation, plan fixed effects, and year fixed effects. The coefficient estimates in panels (b) and (c) are scaled by the standard deviation of the corresponding volatility measure. Annual income buckets on the horizontal axis are $10,000 wide and are based on income values converted to real 2022 dollars. Standard errors are clustered by plan.

Figure A.6: The effect of income volatility on post-separation cashout probability, including separators older than 55



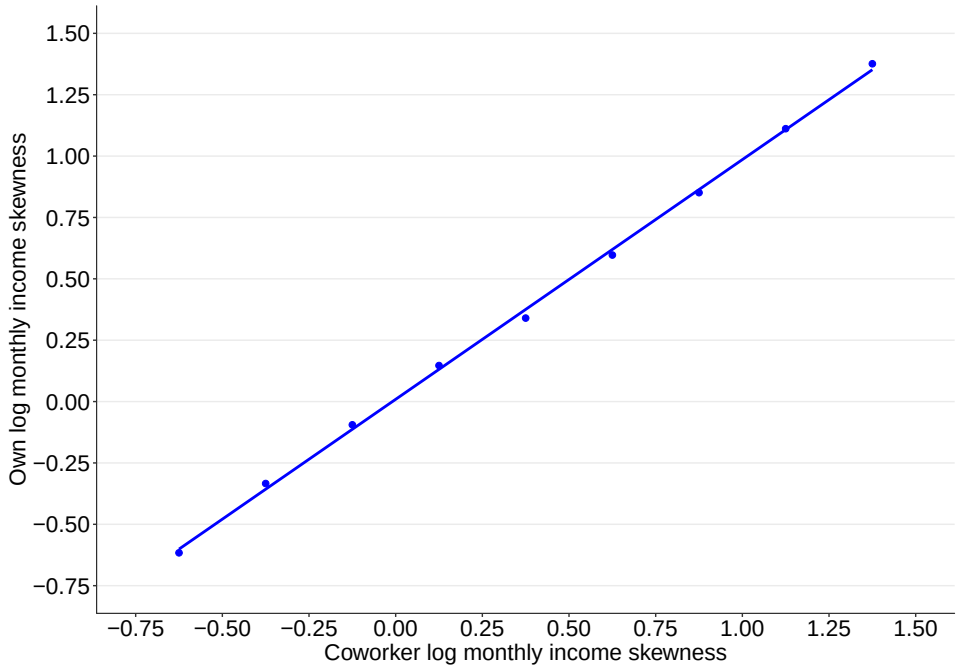
Note: Observations are participants and the sample is restricted to participants who separate from employment during the sample period. Unlike the baseline analysis in Figure 4, I consider all separators, even those who are 55 or older at the end of the calendar year in which they separate and for whom cashouts are not penalized. The income-volatility measures are as defined in Section 3.4 and are averaged across all observable years for a given participant. Participants are classified as hourly if at least 45% of their observable paychecks change by at least $1 relative to the lagged paycheck value. The remaining participants where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. The underlying regression specification in (11) controls for average annual income during the employment spell, age at separation, plan tenure at separation, plan fixed effects, and year fixed effects. The coefficient estimates in panels (b) and (c) are scaled by the standard deviation of the corresponding volatility measure. Annual income buckets on the horizontal axis are $10,000 wide and are based on income values converted to real 2022 dollars. Standard errors are clustered by plan.

Figure A.7: The effect of income volatility on post-separation cashouts, with share of balance cashed out as the dependent variable



Note: Observations are participants and the sample is restricted to participants who separate from employment during the sample period. The dependent variable is the post-separation cashout amount, expressed as a share of the participant's final balance. If the participant does not take a post-separation cash distribution, the dependent variable is 0. The baseline analysis in Figure 4 uses a binary indicator (equal to 1 if the participant takes a post-separation cash distribution and 0 otherwise) as the dependent variable. The income-volatility measures are as defined in Section 3.4 and are averaged across all observable years for a given participant. Participants are classified as hourly if at least 45% of their observable paychecks change by at least $1 relative to the lagged paycheck value. The remaining participants where fewer than 45% of paychecks change by at least $1 are classified as salaried. See Section 3.2 for a discussion of this classification rule. The underlying regression specification in (11) controls for average annual income during the employment spell, age at separation, plan tenure at separation, plan fixed effects, and year fixed effects. The coefficient estimates in panels (b) and (c) are scaled by the standard deviation of the corresponding volatility measure. Annual income buckets on the horizontal axis are $10,000 wide and are based on income values converted to real 2022 dollars. Standard errors are clustered by plan.

Figure A.8: First-stage relationship between own and coworkers' monthly income skewness



Note: This binscatter plot only considers participant-years that are classified as hourly (i.e., participant-years where at least 45% of paychecks change by at least $1 relative to the lagged paycheck value; see Section 3.2 for a discussion of this classification rule. The volatility measure considered here is the skewness of log monthly earnings per paycheck; see Section 3.2 for a definition. I group participant-years into 0.25-unit-wide bins based on average coworker skewness, then plot the average value of own skewness within each bin. A simple bivariate regression of own on coworker skewness yields a slope coefficient of 0.99 with a standard error of 0.0008 and an R-squared value of 0.12.

Appendix B Bonus-identification algorithm

To identify annual bonuses among salaried participants, I use the following algorithm. I only apply the algorithm to participant-years that are classified as salaried (i.e., participant-years in which less than 45% of paychecks change by at least $1 relative to the lagged paycheck value; see the discussion in Section 3.2).

- Let g_{it} be annual base income for salaried participant i in year t. The typical paycheck is therefore $\frac{1}{n_{it}} g_{it}$, where n_{it} is the number of paychecks for participant i in year t.
- When an annual bonus is paid, the paycheck becomes $g_{it}\left(\frac{1}{n_{it}} + \theta\right)$, where θ is the bonus amount as a share of base income.
- Proxy for g_{it} with $n_{it} y_{it}^m$, where y_{it}^m is the median paycheck amount for participant i in year t.
- Choose a minimum bonus-detection threshold $\underline{\theta}$. I set $\underline{\theta} = 0.05$, so I aim to detect bonuses that are at least 5% of base income.
- Classify a given paycheck y_{itp} as containing a bonus if:

$$\begin{aligned} y_{itp} &\geq n_{it} y_{it}^m \left(\frac{1}{n_{it}} + \underline{\theta}\right) \\ &= y_{it}^m (1 + n_{it}\underline{\theta}) \end{aligned}$$

- If the above condition is satisfied for a given paycheck y_{itp}, the inferred bonus amount is $y_{itp} - y_{it}^m$ and the inferred base amount is y_{it}^m.

Table B.1 shows summary statistics for the identified bonuses. About 35% of salaried participant-years have bonus income. In the vast majority of these cases, the bonus is identified in only one paycheck, consistent with the algorithm's goal of identifying annual bonuses. As a share of base income, bonus amounts range from 5.8% at the 10th percentile to 25.8% at the 90th percentile.

Table B.1: Inferred bonuses among likely salaried workers

	Likely hourly	Likely salary
Participant-years	6,294,909	3,087,438
Share with inferred bonus	–	34.9%
# of paychecks with bonus income (cond. nonzero)		
Mean	–	1.2
10th	–	1
50th	–	1
90th	–	1
Bonus as share of base income (cond. nonzero)		
Mean	–	14.5%
10th	–	5.8%
50th	–	10.8%
90th	–	25.8%

Note: Observations are participant-years. Participant-years are classified as hourly if at least 45% of the participant's paychecks in the relevant year change by at least $1 relative to the lagged paycheck value. The remaining participant-years where fewer than 45% of paychecks change by at least $1 are classified as salaried.