FinTech to the (Worker) Rescue? Earned Wage Access, Worker Welfare and Employee Retention*

Jose Murillo Primero AI

Boris Vallee INSEAD

Dolly Yu† Yale School of Management

August 2026

Abstract

We study the usage and welfare implications of earned wage access (EWA) using administrative and survey data from a Mexican FinTech provider. Adoption is significant, and usage concentrates at the end of the pay cycle. We develop a stationary buffer-stock framework, calibrated to the data, in which EWA provides liquidity insurance and consumption-timing alignment. Access is worth about 2.5% of a paycheck per year for the average user, and an order of magnitude more for impatient, low-liquidity workers. Automatic repayment makes these gains more robust to present bias than with revolving credit. EWA usage is associated with higher employee retention, as the framework implies.

Keywords: Earned Wage Access, Consumption Smoothing, Employee Retention, FinTech

JEL classification: D14, D15, G21, G51, J33, O16

*We are grateful to Alex Barbu, John Campbell, James Choi, Claire Celerier, Paul Fontanier, Andreas Fuster (discussant), Paul Goldsmith-Pinkham, Sean Higgins, Ankit Kalda (discussant), Thorsten Martin, Daniel Schmidt, Jialan Wang, Daojing Zhai and Steve Zeldes for useful feedback and thank conference participants at the 2023 NBER Innovative Data in Household Finance conference, the 2025 Paris Finance Meeting, the 2026 Midwest Finance Association Annual Meeting, and the 2026 Swiss Society for Financial Research Conference, as well as seminar participants at Harvard Business School, Yale University and Peking University for their comments and suggestions. We are indebted to Minu for generously providing us access to their data. All errors are ours only.

†Jose Murillo is at Primero AI. Email: jose@primero.ai. Boris Vallee is at INSEAD. Email: boris.vallee@insead.edu. Dolly Yu (Contact Author) is at Yale School of Management. Email: dolly.yu@yale.edu.

1 Introduction

A large share of households across both developed and emerging economies lack the liquidity to absorb a modest expense shock. According to the Federal Reserve's 2023 survey on the economic well-being of US households, 37% of US households could not cover a $400 emergency expense using cash or its equivalent; the fractions unable to raise emergency funds are larger across emerging markets, where formal credit is scarcer and public safety nets are thinner: in the World Bank's Global Findex survey, 45% of adults in developing economies report that they could not come up with emergency funds within a month, and a majority could not do so on short notice (Demirgüç-Kunt et al., 2022). Such cash shortfalls affect households along several dimensions, including consumption (Shapiro, 2005), psychological well-being (Haushofer and Fehr, 2014), and productivity at work (Kaur, Mullainathan, Oh, and Schilbach, 2025). The financial health of employees has moreover been shown to deteriorate before payday (Carvalho et al., 2016).

A promising avenue for improving employee financial health and welfare is therefore for employers to offer benefits that address worker financial strain, as a complement or lower-cost substitute to raising wages. An emerging tool for that purpose is earned wage access (EWA), a benefit facilitated by specialized FinTech firms that allows workers to withdraw at will a fraction of the wages they have earned but have not yet been paid. Proponents argue that the service provides the short-term relief payday lending offers, without the latter's drawbacks, such as predatory pricing and rollovers. The service has spread rapidly across both developed and emerging economies, typically adopted by large employers of low-wage, high-turnover workforces: examples range from Walmart, McDonald's, and Uber in the United States to major retailers, manufacturers, and service firms throughout Latin America, Africa, and Asia. Usage has grown correspondingly: in the United States, employer-partnered providers advanced $22.8 billion across 214 million transactions to an estimated 7.2 million workers in 2022, nearly double the 2020 volume, and, combined with direct-to-consumer providers, the market reached roughly $32 billion (Consumer Financial Protection Bureau, 2024). Adoption has expanded fastest in emerging markets, where alternative sources of short-term liquidity are most limited.

This development raises the question of the economic value of earned wage access. What are the adoption and usage patterns of such a product, and what do they tell us about the

demand for the service? Can it raise worker welfare, and if so, for which workers? Which economic mechanisms explain such an effect? How does earned wage access compare to the alternatives a worker already has, such as conventional credit products or simply being paid more frequently? And are there unintended costs to offering or using it?

This paper tackles these questions in three steps. We first investigate empirically the usage patterns of earned wage access and the motives for its use. We then develop a quantitative framework that pins down the mechanisms underlying the economic value of the product, and, by calibrating it to the usage moments we document empirically, quantifies that value and characterizes the contexts in which it is largest. Last, we relate EWA usage to employee retention and ask whether the resulting patterns are consistent with the framework's welfare implications, in particular whether the retention relationship is concentrated among the workers for whom the framework implies the largest gains. For our empirical analysis, we exploit a novel data set from a leading provider of EWA in Mexico, Minu. This empirical setting is particularly suited to answering our research questions as we are able to access granular usage data, as well as worker demographics and employment history for a large sample of workers, which we complement with survey data from EWA users on their motives for using the product and the financing alternatives it displaces. The context of a developing economy is of particular interest for investigating the benefits of such an offering, as alternative consumption-smoothing tools are expensive and safety nets less protective.

Our empirical results are as follows. We first document a series of stylized facts about the adoption and usage of this novel financial product, typically offered via a partnership between a FinTech firm and a large employer. A significant fraction of workers adopt the service within a short time frame. EWA appears to be more popular with workers who are male and young; in contrast with the widespread view that this product caters mostly to the lowest-income workers, adoption is spread across skill and wage levels. Usage is concentrated toward the end of the pay cycle, when the previous paycheck is more likely to be depleted, and withdrawals are economically significant, averaging 12.1% of the pay-cycle wage for semi-monthly paid workers, while remaining well below the contractual cap of half of the wages earned but not yet paid.[1] Usage patterns are moreover bimodal: the majority of workers use the service occasionally, with withdrawals stable in size and

[1]For weekly paid worker, it represent 29.9% on average.

number across pay cycles, while a subgroup of frequent users relies on the service in most pay cycles, withdrawing a larger share of their paycheck and drawing somewhat earlier within the cycle. Survey responses collected by Minu from 1,435 of its users complement these administrative records. Workers report using EWA for expected and unexpected expenses in roughly equal proportions. A majority of users also reduce their reliance on other financial products once EWA becomes available, with the sharpest substitution away from informal borrowing from family and friends and from pawnshops, two sources likely to carry high costs, whether pecuniary or non-pecuniary.

We then assess these facts against three candidate drivers of the demand for EWA: insurance against unexpected expenditure shocks, alignment of the timing of income with that of predictable expenses, and a preference for immediate consumption stemming from present bias. The evidence supports liquidity insurance and timing alignment as the dominant motives for the majority of users: usage is irregular in timing and amount but tilted toward the end of the pay cycle, and rarely reaches the withdrawal cap, and EWA most strongly displaces the informal borrowing on which constrained workers rely for emergencies. Withdrawals around Semana Santa, a major Mexican holiday that in 2021 fell just before a pay date, speak to the timing-alignment channel: aggregate withdrawals rise sharply in the days preceding the holiday. Since a worker with an adequate buffer would save ahead of such a foreseeable expense, this spike also reveals the presence of liquidity-constrained or present-biased users. Last, the behavior of frequent users, who include higher-rank workers less likely to be constrained, is consistent with a moderate degree of present bias, whereas an extreme version, under which workers unaware of their bias would withdraw the maximum amount as early as possible in the pay cycle, is rejected by both the usage and the survey data. These patterns pin down how workers use EWA but not what it is worth to them. Withdrawals that relieve a binding liquidity constraint and withdrawals driven by present bias may look alike in the data, and usage alone cannot reveal the size of the gain from access, for whom it is largest, or how it compares with the credit these workers would otherwise turn to. Answering these questions requires a model of consumption and precautionary saving in which preferences and access to credit are made explicit, which is what we turn to next.

We thus develop a stationary buffer-stock framework that combines these ingredients to pin down and quantify the economic value of EWA. In the model, a worker paid semi-

monthly faces mid-cycle expenditure shocks and may exhibit present bias. Her precautionary buffer, and hence her exposure to shocks, is an equilibrium object. The framework rationalizes the two usage patterns we document, occasional shock-driven use and regular use by a subgroup of frequent users, as the equilibrium behavior of time-consistent and present-biased buffer-stock savers, respectively, and delivers a sharp characterization: workers with interior savings use EWA only against shocks, while hand-to-mouth workers also use it to align consumption with the timing of pay.

The framework first establishes for whom the product is most valuable: impatient, illiquid, risk-averse workers facing sizable expenditure shocks, which corresponds to the profile of the workforces to which EWA is marketed. Across these dimensions, the value of access varies by an order of magnitude. Present bias operates through two opposing channels: it thins precautionary buffers, which increases exposure to shocks and raises the value of the option, but the resulting heavier usage offsets part of that value in fees. For sizable shocks the exposure channel dominates, and present-biased workers value the product most. They are therefore also its most intensive users, a pattern that raises the concern that the product profits from their bias. Our welfare results indicate the opposite: because repayment is automatic and tied to the next paycheck, heavier usage cannot compound into debt, and the present-biased worker retains almost the same gain from access as the time-consistent one. Revolving credit, which leaves the timing of repayment to the worker, offers no such discipline.

A comparison with traditional credit products isolates the role of the product's automatic repayment feature, and of its lower cost resulting from the absence of credit risk. Workers who are aware of their own bias would repay anyway, so the feature costs them little. For workers naive about their own bias, in contrast, revolving credit generates compounding rollover debt and can destroy welfare, while the harm from EWA is bounded by its fees, as tying repayment to the paycheck disciplines borrowing. Calibrated to our usage data, the model implies that EWA access is worth about a tenth of a semimonthly paycheck in present value for the average user, a lower bound on its potential for more constrained populations, and close to the value of a frictionlessly priced credit line, while the payday-style credit actually available to this population delivers only about a tenth of that value, at two and a half times the cost per unit of liquidity. For the present-biased type, willingness to pay barely exceeds the long-run value of access, and the harm from overuse

stays bounded as bias deepens. We further compare EWA with more frequent pay, which targets the same underlying friction: while raising pay frequency can substitute for part of EWA's role, it is a far costlier and often infeasible change for the employer, whereas EWA delivers state-contingent liquidity on demand at low cost and layers onto any existing pay schedule. Finally, we derive several insights for EWA product design from these results.

We finally confront the framework's central implication, that EWA generates surplus concentrated among constrained workers, with data on employee retention. Workers using EWA have a significantly lower probability of leaving their firm in any period than workers who do not, controlling for age, gender, starting month, and job position. The association is concentrated among the regular users whose behavior the framework reads as hand-to-mouth, which is where it implies the largest gains, and is of similar size across job ranks. We estimate these relationships with a proportional hazard model on a worker-level panel, and they are robust to a discrete-time hazard specification and to a propensity-based difference-in-differences design. These are associations rather than causal effects: adoption is a worker's decision, and our designs mitigate but do not eliminate selection on unobserved attachment to the firm.

Our study relates to several streams of literature. First, our study contributes to the literature on cash shortfalls and their effects (see for instance Shapiro, 2005; Fink et al., 2020). A large body of work documents how adverse expense shocks, often health-related, impose sizable and persistent financial costs on households (Dobkin et al., 2018, Kluender et al., 2025). By studying a tool that aims to alleviate such shortfalls, we provide evidence on the value of insuring against them. Our work also connects with the literature on short-term credit, which finds that payday lending can play a liquidity insurance role (Morse, 2011) but often carries large unintended consequences, including persistent defaults and repeat borrowing (Melzer, 2011, Gathergood et al., 2019).

Second, this study speaks to the effects of FinTech development on household welfare. Technological change has reshaped the supply of consumer credit (Berg et al., 2022), often by exploiting new information, as when FinTech lenders use borrowers' cashless-payment data to screen and extend credit (Ghosh et al., 2026), much as EWA uses the employer's payroll integration. This literature has documented contrasting effects for household finances, from overborrowing (Di Maggio and Yao, 2021, Chava et al., 2021) to improved asset allocation (Reher and Sokolinski, 2024). A related strand studies new FinTech prod-

ucts delivered through an intermediary, such as buy-now-pay-later offered via merchants (Berg et al., 2025), paralleling earned wage access offered via employers. Earned wage access itself is only beginning to be studied. Indarte et al. (2026) measure and diagnose consumption distortions, using data from an EWA provider to separate liquidity constraints from behavioral ones, a distinction our framework also draws and which their evidence also supports. Our contribution is to price the product: we compute the welfare value of access, the willingness to pay, and the gap between them, and we benchmark EWA against two credit regimes and against a higher pay frequency, which delivers the counterfactuals a provider or an employer actually chooses among.

Third, by studying the mechanisms that rationalize the adoption of this innovation and its effects, our work also speaks to the optimal frequency of pay (Parsons and Van Wesep, 2013, Baugh and Correia, 2022), to which we contribute an explicit welfare comparison between EWA and more frequent pay, as well as the utility function of individuals, such as the well-documented existence of present bias (Laibson, 1997, Parker, 1999, Stephens, 2003, Shapiro, 2005, Olafsson and Pagel, 2018) and its consequences for consumption and welfare (Maxted, 2025), and how such non-standard preferences can influence financial product design (Calvet et al., 2023).

Last, our work adds to the literature on strategies to reduce firm turnover and foster retention, which include pay design (Oyer, 2004, Oyer and Schaefer, 2005), human resource management (Burks et al., 2015) or corporate culture, including CSR (see for instance Carnahan et al., 2017, Bode et al., 2015). Most directly, higher pay itself reduces worker turnover and raises productivity (Dube et al., 2016, Emanuel and Harrington, 2024), but raising wages is costly for the firm. Non-pay amenities can achieve similar retention gains at lower cost: Bloom et al. (2024) show in a randomized experiment that granting hybrid work from home cuts quit rates by a third without hurting performance. Earned wage access is such an amenity, a cost-effective tool that can substitute for or complement raising pay and other corporate actions targeting retention.

The remainder of the paper is structured as follows. Section 2 provides background on earned wage access and describes our data. Section 3 documents stylized facts about earned wage access usage. Section 4 assembles direct evidence on the drivers of the demand for EWA from usage diagnostics and survey data. Section 5 develops a stationary framework, disciplined by our usage data, to pin down and quantify the economic value of EWA. Section

6 relates earned wage access usage to employee retention and asks whether the patterns are consistent with the framework's welfare implications. Section 7 concludes.

2 Background and Data

2.1 Background on Access to Earned Wages

Earned wage access services, also referred to as on-demand pay, allow users to withdraw wages earned during the pay cycle before the actual payday. The service is typically offered in one of two ways by a FinTech firm: either it partners with employers and integrates the product into the employer's payroll system, serving as a technology provider, or it acts as an intermediary: it obtains information from the employer on the level of earned wages, transfers cash directly to workers' bank accounts, and is reimbursed by the employer on payday. While the two arrangements imply different business models for the FinTech firm, they deliver a similar product to the end user. Over the past decade, EWA has become an increasingly popular benefit offered by employers for worker retention, both in developed economies (with adopters such as Walmart, Amazon, and McDonald's in the United States) and across emerging markets, typically among large employers of low-wage service workers. In the United States alone, employer-partnered providers advanced $22.8 billion in 2022, nearly double the 2020 volume (Consumer Financial Protection Bureau, 2024). A 2019 Visa survey reports that 95% of workers would be interested in working for companies providing EWA (Visa, 2019). Earned wage access is not considered a loan in most jurisdictions, and is therefore legally distinct from payday lending. Earned wage access has nonetheless raised consumer-protection discussions, mostly in cases where withdrawal fees correspond to high annualized rates.

The employer-partnered earned wage access we study should be distinguished from direct-to-consumer cash advance products, which are sometimes marketed under similar labels. Employer-partnered EWA is offered as a benefit through the employer, is tied to wages the worker has already earned as recorded by the payroll system, and is repaid automatically by deduction from the next paycheck; it therefore bears essentially no credit risk and carries low, often employer-subsidized, fees. Direct-to-consumer cash advances, by contrast, are marketed directly to consumers with no employer relationship, estimate income from bank-account records rather than verified earned wages, and are repaid by

debiting the user's bank account; they bear default risk, frequently rely on tips and subscription fees, and are conceptually closer to payday lending, the key distinction being that no rollover is possible. Our data and analysis concern the employer-partnered model.

2.2 Data

We obtain our data from Minu, a leading provider of EWA in Mexico. Minu follows the intermediated model and provides a mobile application that allows the employees of its corporate clients to transfer earned but not yet paid wages to their bank account at any point within the pay cycle. The service is offered via the employer as a benefit to its workers. Workers can transfer any amount up to a cap, any number of times, during the pay cycle. The next paycheck is automatically lowered by the total amount withdrawn during the past pay cycle. No cost accrues beyond the fixed fee, so the paycheck adjustment corresponds exactly to the cumulative amount withdrawn. As of 2022, the amount withdrawn is subject to a cap of 50 percent of available earned wages at the time of withdrawal, and each transfer carries a fixed service fee of 39 pesos (around US$2). Who bears this fee is a decision of the employer: providers offer plans under which the firm absorbs it as part of the benefit, so that the worker withdraws at no cost. On the average withdrawal in our data, about 783 pesos, the 39-peso fee represents close to 5%, which, given that advances are repaid within the pay cycle, corresponds to a high annualized rate, above the roughly 110% illustrative APR that Consumer Financial Protection Bureau (2024) reports for the average US employer-partnered earned-wage transaction. Where the worker bears it, EWA is therefore not an inexpensive form of short-term liquidity, although it remains well below the annualized cost of payday lending, and the employer can absorb the fee altogether.

We obtain data covering the employees of two corporate clients of Minu, both service firms. For each company, we obtain two sets of data. The first contains user activity from the Minu app: transaction-level details of all withdrawals, including time and date, amount, and fees, for all employees registered to use Minu. The coverage window differs across the two firms: for one, these app data run from January to July 2020; for the other, from September 2020 to April 2021. The second is shared by Minu's corporate clients and includes basic demographic characteristics and work histories for all of their employees, regardless of Minu usage. In particular, we obtain the start date, end date, and gender of all employees. For a subset of employees, we also know their birth year, location of

work, current and previous job positions and rank. These employee records cover the workers' full employment spells, with hire dates predating each firm's withdrawal window by several years and separations observed through the end of that window (July 2020 and April 2021, respectively). We combine all of those types of information and construct a panel at the employee–pay cycle level. Two features of this sample bound what it can speak to. It covers one provider and two employers, both large service firms with low-wage, high-turnover workforces of the kind to which EWA is marketed, so our estimates describe that setting rather than the Mexican workforce as a whole. Its withdrawal windows also fall in 2020 and 2021, a period of unusual labor-market disruption, which we take into account when interpreting the separation patterns of Section 6.

To further investigate the mechanism underlying the demand for access to earned wages and its association with higher retention, we collaborated with Minu to collect survey data from its EWA users in July 2023. Minu sent a 5-question survey through WhatsApp to 9,700 users in around 20 companies to better understand users' purposes and usage behavior. We received 1,435 responses within two weeks of the survey launch, a 15% response rate. The survey asks about the main purposes of use, the alternative products used before and after EWA adoption, typical withdrawal amounts, and a comparison with other company benefits.

2.3 Summary Statistics

There are a total of 44,041 workers in our sample, and about 9% of them use EWA during their employment. Panel A of Table 1 presents demographic characteristics of EWA users versus non-users. The two groups have similar age distributions, with a mean age of about 33, and about seven in ten workers in each group are paid semimonthly rather than weekly. EWA users are more likely to be male, started at the company earlier and therefore have somewhat longer tenure, and are higher ranked than non-users. Panel B presents statistics on EWA usage, separately for semimonthly- and weekly-paid users. In both groups most users make fewer than two withdrawals per pay cycle. The average amount withdrawn per pay cycle is 953 pesos for semimonthly-paid users and 385 pesos for weekly-paid users; since the two groups draw very different paychecks, the comparable statistic is the share of the paycheck withdrawn, which is 12% for the semimonthly-paid and 30% for the weekly-paid. Semimonthly-paid users are also active in a much larger share of their own pay cycles, 58%

against 32%.

[INSERT TABLE 1 HERE]

3 Stylized Facts about Earned Wage Access Usage

In this section, we present novel stylized facts about earned wage access adoption and usage. We organize them into four sets: which workers adopt the product; how adopters use it, namely the frequency, size, and timing of withdrawals; the heterogeneity between occasional and frequent users; and survey evidence on the purposes of use and the products EWA replaces. Beyond documenting how this novel product is used, these facts play two roles in the rest of the paper: they provide the diagnostics with which we assess the drivers of demand in Section 4, and the moments to which the quantitative framework of Section 5 is calibrated.

3.1 Worker Characteristics and EWA Adoption

First, we document that when an employer offers EWA to its employees, a significant share of workers adopt the service within a short time frame. As shown in Panel A of Table 1, about 9% of workers from our sample made withdrawals through EWA at least once after the introduction of the service. The share of workers using Minu also trends upward in the pay cycles following the introduction, suggesting growing popularity as workers become aware of the service.

We then explore which worker characteristics predict the use of earned wage access. Using a logistic model, we regress an indicator for using EWA on worker characteristics, and report the coefficients in Table 2. Male and younger workers are significantly more likely to use the service, and the implied magnitudes are large: relative to the 7% baseline adoption rate in the estimation sample, being male raises the likelihood of using EWA by about half of that baseline and being young by about a third. EWA also appears to be less popular with employees of lower rank, and the gradient is steep: the shares in Panel A of Table 1 imply adoption rates of about 8%, 14% and 23% for low-, mid- and high-rank workers. Who takes up the product is not who values it: a new benefit delivered through an application is adopted first by the workers most equipped to use one, as is typical of technology diffusion. Conditional on adopting, usage runs the other way, as low-rank users

withdraw a larger share of their wages (Appendix Figure A.3), and the welfare numbers of Section 5 accordingly describe the average user rather than the average worker. Last, starting after the introduction of EWA predicts significantly lower adoption, which could reflect both lower exposure to communication about the benefit and less time to use the service. Figure A.1 plots the predicted probability of EWA usage across worker subgroups implied by these estimates.

[INSERT TABLE 2 HERE]

3.2 EWA Usage Patterns

We now zoom in on the workers who use earned wage access and investigate their usage patterns. We first explore the distribution of the frequency, amount, and timing of withdrawals, and how these evolve over the seven months following first adoption. Figure 1 presents the distribution of withdrawal timing (Panel A) and amount, scaled by 2-week pay cycle income (Panel B) and earned wage at the time of withdrawal (Panel C), all at the worker–pay cycle level and conditional on EWA usage by that user in that pay cycle.

We observe from Panel A that the majority of users use the service in the second half of the pay cycle, which corresponds to the period in which the worker has accumulated earned but unpaid wages, and is also more likely to have exhausted the previous paycheck. Panel B shows that workers withdraw economically significant amounts: the average amount withdrawn over a pay cycle equals 12.1% of the cycle wage. At Minu, the amount workers can withdraw is capped at 50% of the wages already earned but not yet paid. Panel C shows that most workers withdraw well below this cap, even though the fixed per-transfer fee makes larger withdrawals cheaper per peso withdrawn.

[INSERT FIGURE 1 HERE]

Figure 2 documents how withdrawal activity evolves over pay cycles for an average user. After their first use, around 40% of EWA users keep using the service in every period. Moreover, conditional on withdrawing in a given pay cycle, users withdraw increasingly large amounts: the average user's withdrawals rise from 11% to 14% of the pay-cycle wage within six months of adoption. We also observe a shift toward earlier withdrawals within the pay cycle, and a slight increase in the number of withdrawals per cycle.

[INSERT FIGURE 2 HERE]

We next explore the heterogeneity of usage patterns across groups of EWA users. In Appendix Figure A.3, we split workers into low- and high-skill positions and plot the same graphs as in Figure 2. Overall, we observe comparable patterns between the two groups, with low-skill workers withdrawing a larger share of their wages, and doing so in fewer withdrawals. These differences are mostly in levels: the two groups exhibit comparable evolutions of usage since their first use.

About three in ten users are paid weekly rather than semimonthly, and they too use EWA regularly, in 31.8% of their own pay cycles against 57.8% for semimonthly-paid users (Panel B of Table 1). Because their cycles are half as long, this amounts to slightly more withdrawal episodes per month. Conditional on use they withdraw a much larger share of their paycheck, 29.9% against 12.1%, and bunch against the cap of 50% of earned wages (Panel C of Figure 1), where semimonthly-paid users remain well below it. This bunching is largely mechanical: at any point in the cycle, a weekly-paid worker has accrued far fewer wages in pesos, so the cap binds at a much smaller withdrawal. Section 5.8 relates these differences to the shorter cycle and to the tighter constraints of this group.

3.3 Frequent Users

Panel A of Figure 2 reveals a subgroup of users who keep using the service consistently after a few pay cycles. We therefore split users into those who use EWA in five or fewer pay cycles over our sample period and those who use it in more than 5 times, and compare their usage patterns.

Figure A.5 in the Appendix plots the same graphs as Figure 2 splitting by usage frequency. Withdrawal amounts as a share of income are similar at first use, after which frequent users settle at a higher level, between 12% and 15% of the paycheck, against 9% to 13% for the other users. Frequent users also draw earlier within the pay cycle throughout, around the middle of the cycle rather than three-fifths of the way through it, although both groups shift somewhat earlier over time. Last, frequent users consistently make more withdrawals per pay cycle than their counterparts.

In summary, usage patterns are bimodal. The majority of workers use the service five times or fewer in the seven months following adoption, with withdrawals that are stable

in size and number across cycles. The remaining users rely on Minu consistently, in most pay cycles, withdrawing larger amounts, more often, and earlier within the cycle. This bimodal pattern, together with the moments documented above (usage in 48% of user–pay cycles, withdrawals averaging 12.1% of the cycle wage conditional on use, and rarely at the cap), constitutes the empirical target of the framework of Section 5, where the two groups emerge as the equilibrium behavior of buffer-stock workers without and with present bias.

3.4 Survey Evidence on Purposes of Use

Survey evidence collected from users complements the administrative records and speaks to why workers use EWA and which products it replaces. These responses are self-reported and cover a subset of users, so we read them as suggestive rather than as establishing the patterns below. Workers report unexpected expenses, namely unexpected bills and medical emergencies, and expected expenses with roughly equal frequency (Figure 3). The size of withdrawals appears to track this split: users report withdrawing below the maximum mostly for daily and expected expenses, and withdrawing the maximum mostly when facing an unexpected expenditure (Figure A.6). The responses also suggest that EWA substitutes for other sources of liquidity, and does so unevenly (Figure 4). Before adoption, most workers report first turning to family and friends or their own savings, then to credit cards or other financial institutions, and rarely to employer loans or other non-financial products. After adoption, reported informal borrowing from family and friends and from pawnshops falls the most, whereas the use of credit cards and instant loans falls the least, a pattern consistent with EWA displacing first the sources of liquidity that are costliest, whether in fees or in the non-pecuniary cost of asking. Reported use of store credit also falls by about two thirds, from 10% to 3% of respondents, a substitution that matters because the store cards serving this population are among the most expensive forms of credit in Mexico, with total annual costs above 80% at the largest chains. Workers also rank EWA as the second most preferred company benefit, behind only medical insurance (Figure A.7).

[INSERT FIGURES 3 and 4 HERE]

Having documented how workers adopt and use EWA and what they report about its purposes, we next confront these facts with the candidate drivers of demand.

4 Evidence on the Drivers of EWA Demand

The facts of the previous sections raise the question of what drives worker demand for earned wage access. We consider three candidate explanations: insurance against unexpected expenditure shocks; alignment of the timing of predictable expenses with that of pay; and present bias, under which workers derive higher utility from consuming earlier within the pay cycle even absent any liquidity need.[2] These drivers are not mutually exclusive, and may coexist across and even within users. Rather than running a horse race among them, this section confronts their distinctive predictions with our usage data and with survey evidence collected from users, and establishes that each explanation has some empirical support. The framework of Section 5 then embeds all three, lets equilibrium behavior determine their respective roles, and quantifies their welfare implications for workers.

Liquidity insurance Cash shortfalls caused by unexpected expenses, such as a health-care bill or an urgent repair, are pervasive among low-wage workers, and their detrimental effects on decision-making and productivity are well documented (Shapiro, 2005, Haushofer and Fehr, 2014, Kaur et al., 2025). If EWA insures such shocks, usage should be irregular in timing and amount, concentrated toward the end of the pay cycle when the previous paycheck is depleted, more valuable to financially constrained workers, and should crowd out costlier sources of emergency liquidity. Both bodies of evidence support this account for the majority of users. In the usage data, withdrawals are irregular across pay cycles, concentrated in the second week of the cycle, and rarely reach the contractual cap. In the survey, unexpected bills and medical emergencies account for roughly half of reported purposes (Figure 3), the cap is reached primarily in response to such unexpected expenditures (Figure A.6), and EWA most strongly displaces the informal and pawnshop borrowing on which constrained workers rely for emergencies (Figure 4). If EWA provides such insurance, its benefits should moreover be concentrated among lower-rank workers, who are the most likely to be constrained.

[2] As explained in O'Donoghue and Rabin (2015), the noticeable discounting that individuals display over horizons of days or weeks would, under time-consistent exponential discounting, imply implausibly extreme discounting over long horizons; such short-run impatience is therefore evidence of present bias.

Timing alignment Households experience shortfalls even absent uncertainty, when predictable expenses fall just before payday (Baugh et al., 2018). EWA alleviates such systematic shortfalls by aligning the timing of income with that of expenditures. The survey points to the quantitative importance of this motive: expected and daily expenses account for roughly half of reported purposes (Figure 3), and it is for such expenses that users tend to withdraw below the maximum (Figure A.6). The usage data isolate the driver directly. Figure 5 exploits the celebration of Semana Santa, whose expenditures are largely predictable and which in 2021 fell at the end of March, just before a pay date: EWA usage rose sharply in the corresponding days. Through the lens of the model, this spike is itself informative about who uses the product. A time-consistent worker with an interior buffer would anticipate a predictable expense and save ahead of it. The Semana Santa spike is therefore suggestive of the existence of a hand to mouth population that taps into EWA to finance expenditures that could have been anticipated.

[INSERT FIGURE 5 HERE]

Present bias In its extreme form, in which strongly biased workers are unaware of their bias, as in the payday-lending analysis of Parsons and Van Wesep (2013), present bias predicts regular and predictable usage, at the maximum the contract allows and unrelated to financial constraint. The timing prediction has to be stated relative to the contract: since wages accrue over the cycle and the cap is set at half of the wages already earned, nothing is available on payday and drawing earlier means drawing less. An unaware present-biased worker would therefore not withdraw on the first day of the cycle, but would withdraw in essentially every cycle, and at the cap given the wages accrued at the time of her request.[3] Both sources of evidence reject this extreme version: in the usage data only a small fraction of withdrawals reach the cap, and utilization of accrued wages stays well below it (Panel C of Figure 1), while in the survey the maximum is reached mainly for unexpected shortfalls rather than as a matter of course (Figure A.6). A subgroup of users, however, exhibits a moderate version of this pattern: the frequent users of Section 3 use the service in most pay cycles, withdraw a larger share of their paycheck and draw somewhat earlier within it

[3]Earlier studies have shown that present bias is more frequent among low wage workers (Meier and Sprenger, 2015), but recent evidence (Olafsson and Pagel, 2018) documents significant present-bias behavior among mid to high income workers.

than other users, though still around mid-cycle and well below the cap, and include higher-rank workers who are less likely to be financially constrained. The Semana Santa spike documented above points in the same direction. Drawing against a foreseeable expense rather than saving for it arises at hand-to-mouth states, which a worker reaches either because she is liquidity-constrained or because she is present-biased, and the usage data alone cannot say which. Separating the two, and pricing what each implies for welfare, is one of the tasks we assign to the framework of Section 5, where the two mechanisms generate the same behavior in the data but different welfare consequences.

Taken together, the evidence indicates overlapping drivers with clear heterogeneity across users. The purposes reported in the survey, unexpected and expected expenses in roughly equal shares, and the substitution away from informal borrowing point to liquidity insurance and timing alignment as the dominant motives for the majority of users; the behavior of frequent users is consistent with a moderate degree of present bias; and extreme, unaware present bias is rejected. These are the ingredients of the framework we develop next: mid-cycle expenditure shocks generate irregular, insurance-driven usage; binding liquidity constraints generate timing-driven usage even absent shocks; and quasi-hyperbolic preferences generate the frequent-user pattern. In equilibrium, the framework reproduces these patterns, including, as an untargeted moment, the survey's roughly equal split between unexpected and expected purposes, and quantifies the welfare gains that each driver generates.

5 A Quantitative Framework for Earned Wage Access

Our empirical analysis documents fast adoption of earned wage access, two distinct usage patterns among adopters, and an association between usage and lower employee turnover. Yet, without a theoretical framework, these findings remain difficult to interpret in terms of underlying economic forces and worker welfare implications. A framework serves three purposes. First, it pins down the sources of demand for EWA, and clarifies why self-insurance through savings or traditional credit products may fail to replicate its effects. Second, once disciplined by the usage moments we observe, it quantifies the welfare gain that EWA generates, and identifies for whom this gain is largest. Third, it allows us to ask whether the demand for EWA emanating from present-biased workers reflects a welfare

gain or the exploitation of their bias, a question at the center of the regulatory debate on this product.

To this end, we develop a stationary buffer-stock model of a worker who is paid bi-monthly, faces mid-cycle expenditure shocks, and may exhibit present bias. The stationary formulation abstracts from the drift documented in Section 3, whereby withdrawals grow and move earlier in the cycles following adoption, which the framework does not generate and which our observation window is too short to tell apart from a transition toward a steady state. Relative to a finite-horizon setting, it has two key advantages. First, precautionary savings are an equilibrium object rather than a parameter: the worker's liquidity buffer, and hence her exposure to shocks, responds endogenously to the availability of EWA. Second, the worker's valuation of EWA is free of terminal-date effects: every pay cycle is identical ex ante, so the value of the product reflects only its insurance and consumption-timing services, and the comparison with credit products is not contaminated by end-of-horizon borrowing motives.

In the model, workers with an interior savings buffer use EWA only against shocks, while workers who are hand-to-mouth mid-cycle also use it to align the timing of consumption with income. The two empirical usage groups we document, occasional users whose withdrawals are irregular and frequent users who draw regularly and increasingly early, emerge as the equilibrium behavior of, respectively, time-consistent and present-biased buffer-stock savers. We proceed in two steps. We first characterize, in general terms, for whom the product is most valuable, tracing the value of access across worker circumstances (liquidity, patience, and shock exposure) and preferences (present bias and risk aversion). We then calibrate the model to the usage moments of Section 3 to quantify the welfare gain for the average user in our data, relating it to this general characterization, before comparing EWA with credit alternatives and drawing implications for product design.

5.1 Model Setup

Time is infinite and each pay cycle consists of two sub-periods, which we refer to as weeks. In the first week of each cycle (the *payday stage*, denoted A), the worker receives her semimonthly wage, normalized to $w = 1$. In the second week (the *mid-cycle stage*, denoted B), no income arrives and the worker faces an expenditure shock ξ with probability p. The shock captures an unavoidable expense that must be resolved within the period, such as a

medical bill, an urgent repair, or an obligation to family. This timing captures the central friction motivating EWA: pay is intermittent while consumption needs are continuous, and unplanned needs may arrive precisely when the previous paycheck has largely been spent.

The worker holds a liquid balance earning no interest, and cannot borrow in the baseline. At the payday stage, with cash-on-hand m (wage plus any carried savings, net of EWA repayments and access fees defined below), she chooses consumption c_A and carries savings $s = m - c_A \geq 0$ into the mid-cycle stage. At the mid-cycle stage, if she has access to EWA, she can withdraw an amount $x \in \{0\} \cup [x_{\min}, \bar{x}]$ of her earned but unpaid wages, receiving ηx with $\eta < 1$: the wedge $1 - \eta$ is the proportional withdrawal fee and $x_{\min}$ a minimum withdrawal size. The withdrawal is automatically deducted from the next paycheck, the defining institutional feature of the product. The cap $\bar{x}$ reflects the contractual limit on the fraction of earned wages that can be advanced.[4] She then chooses mid-cycle consumption c_B, and carries $k = s - \xi \mathbf{1}_{\text{shock}} + \eta x - c_B \geq 0$ into the next payday, at which point cash-on-hand is $m' = k + w - x - \pi$, where π is a per-cycle access fee that we use as a pricing device below (EWA is in practice offered at $\pi = 0$).[5]

The worker has constant relative risk aversion utility $u(c) = c^{1-\gamma}/(1-\gamma)$ and quasi-hyperbolic preferences (Laibson, 1997): viewed from any week t, lifetime utility is

$$U_t = u(c_t) + \beta \sum_{j \geq 1} \delta^j \, \mathbb{E}_t[u(c_{t+j})],$$

where δ is the weekly exponential discount factor and $\beta \in (0, 1]$ captures present bias, with $\beta = 1$ corresponding to a time-consistent worker. Workers are sophisticated about their own bias, that is, they correctly anticipate the behavior of their future selves; sophistication here

[4] At Minu, withdrawals are capped at 50% of earned but unpaid wages, so the cap in units of the cycle wage depends on how far into the cycle the worker draws. Across the withdrawals in our sample the cap stood at 27% of the worker's own paycheck on average. We set $\bar{x} = 0.30$. The withdrawal fee is fixed (39 pesos per transfer); we approximate it with a 2% proportional cost ($\eta = 0.98$), its approximate burden on the larger withdrawals that account for most of the amounts withdrawn. Modeling the fixed fee explicitly would leave the analysis qualitatively unchanged: the marginal peso withdrawn is repaid one-for-one out of the next paycheck, so it is the repayment, not the fee, that prices the marginal withdrawal, and interior withdrawals survive rather than bunching at the cap. The fee instead operates on the extensive margin, ruling out small withdrawals, and Proposition 1 continues to hold with the fee as the hurdle deterring usage by unconstrained workers. We impose that margin directly through the minimum withdrawal size, setting $x_{\min} = 0.015$ in units of the cycle wage, or about one hundred pesos: no withdrawal smaller than this minimum appears in our data.

[5] We impose a subsistence consumption floor $c_{\min} = 0.02$ to ensure feasibility in extreme states in which the shock exceeds available resources; the floor binds with negligible probability in all equilibria we compute.

refers to self-awareness of the bias and not to financial sophistication in any broader sense.[6] We solve for a stationary Markov-perfect equilibrium of the intra-personal game: each week's self chooses optimally taking the behavior of her future selves as given, and policies at each stage depend only on cash-on-hand. Choices at each stage are subject to small independent taste shocks of scale σ, which capture random variation in the momentary attractiveness of drawing on earned wages, such as a scheduled obligation falling due, so that two workers in the same measured state need not act identically. This is the standard smoothing device for dynamic discrete choice with a continuous margin (Rust, 1987, Iskhakov et al., 2017). We set $\sigma = 3 \times 10^{-4}$, worth about a quarter of a peso at our calibration against a mean withdrawal of around eight hundred pesos, so that it disciplines the numerical solution without bearing on the economics. The numerical procedure is described in Appendix A.

5.2 When is EWA Used? Insurance versus Consumption Timing

The model yields a prediction on the composition of usage, which formalizes the discussion of Section 4.

Proposition 1. *In any stationary equilibrium, consider a mid-cycle state in which no shock has occurred and the worker's savings choice is interior ($k^* > 0$). Then the optimal EWA withdrawal is $x^* = 0$, for any degree of present bias $\beta \in (0, 1]$.*

The proof, in Appendix A, rests on a simple perturbation argument: a worker with interior savings who withdraws could instead reduce both her withdrawal and her savings so as to leave current consumption unchanged, and would arrive at the next payday with strictly more cash, having saved the fee. Intuitively, when savings are interior the worker is already shifting resources across weeks at a zero interest rate, so paying the fee $1 - \eta$ for the same transfer is strictly dominated. Two features of this result matter. First, it holds for present-biased workers as well: present bias makes the worker want to consume more today, but this preference is already reflected in her savings choice; EWA offers no cheaper way to tilt consumption when savings are interior. Second, the argument applies

[6] The survey and usage evidence of Section 4, in which even regular users withdraw well below the cap, is inconsistent with extreme unawareness. We relax sophistication in the comparative statics and credit comparison of Section 5.7, where naivete is precisely what gives the automatic repayment feature of EWA its disciplining value.

identically in shock states, so *any* EWA usage must occur at states where the worker is hand-to-mouth mid-cycle ($k^* = 0$), that is, where she would like to consume more than her remaining liquidity ahead of a paycheck she cannot yet access. At such corner states, EWA relaxes a binding liquidity constraint; usage without a shock then serves purely to align the timing of consumption with that of income.

Equilibrium usage therefore decomposes into usage in shock states, which provides insurance, and usage in no-shock states, which aligns consumption with the timing of pay; the share of each is an equilibrium outcome determined by how thin buffers are. Present bias, by eroding buffers, shifts workers toward the corner region and generates regular and predictable usage, the behavior of the frequent users documented in Section 3, whereas time-consistent workers with adequate buffers use EWA only occasionally, when shocks find them with depleted liquidity. The survey evidence that withdrawals are used about equally for unexpected and expected expenses (Figure 3) maps directly onto this decomposition. Figure 6 illustrates the insurance service directly: when a mid-cycle expenditure shock finds a worker with a depleted buffer, consumption drops sharply without EWA, whereas access to earned wages keeps the consumption path nearly flat, at the cost of a mild drag over the following cycles as the advance is repaid. The effect is more pronounced for the present-biased type, whose thinner buffer leaves her more exposed when the shock hits.

[INSERT FIGURE 6 HERE]

5.3 Pricing EWA: Willingness to Pay versus Welfare Gain

We measure the value of EWA with two per-cycle access prices, which differ only in the welfare criterion, and a state-contingent capitalized value. Let $V(m)$ denote the value, under exponential discounting δ, of the equilibrium consumption process of a worker with payday cash-on-hand m; this *long-run criterion* is standard for the welfare evaluation of present-biased agents (O'Donoghue and Rabin, 2015). Let $\tilde{V}(m)$ denote the corresponding value under the payday self's actual preferences, that is, with future utils weighted by β.

The *welfare-consistent price* π^*_{welf} is the per-cycle access fee that leaves the long-run criterion indifferent between having and not having EWA, evaluated at the ergodic mean cash-on-hand of the no-EWA economy, $\bar{m}$ (the worker's average payday cash-on-hand in

the long run):

$$V^{\text{EWA}}(\bar{m}; \pi^*_{\text{welf}}) = V^{\text{No EWA}}(\bar{m}).$$

It is a flow equivalent variation: the recurring fee, in units of the semimonthly wage, whose payment would exactly exhaust the welfare gain from the product. The *willingness to pay* π^*_{WTP} is defined identically but using $\tilde{V}$, and captures the participation margin: the maximum recurring fee at which the worker's deciding self enrolls. For time-consistent workers the two coincide; for present-biased workers the difference $\pi^*_{\text{WTP}} - \pi^*_{\text{welf}}$, which we label the *internality gap*, measures the extent to which the product's appeal exceeds its long-run value. This gap is the model counterpart of the concern that EWA might monetize present bias.[7]

Finally, to characterize how the value of EWA varies with a worker's financial position, we compute the one-time payment $\tau(m)$ that makes a worker without EWA as well off as one with permanent free access, state by state: $V^{\text{No EWA}}(m + \tau(m)) = V^{\text{EWA}}(m; 0)$. The function $\tau(\cdot)$ capitalizes the flow value of access and reveals its concentration among low-liquidity workers.

5.4 For Whom is EWA Valuable?

Before taking the model to our data, we characterize how the value of EWA access varies with the worker's circumstances and preferences. Figure 7 traces the pricing objects defined above across the parameter space, varying one dimension at a time around a benchmark parameterization (that of the calibration in the next subsection).

Liquidity The value of access is strongly decreasing in the worker's financial slack (Panel A). The capitalized value $\tau(m)$ exceeds 20% of a paycheck for a worker entering payday with half a paycheck of resources, and declines steadily as cash-on-hand improves: EWA is an option on liquidity, and like any option it is most valuable close to where it binds.

[7] Because the equilibrium at any fee is solved under the worker's actual (quasi-hyperbolic) behavior, both prices reflect the same usage; only the evaluation criterion differs. Neither price is an actuarially fair premium. The provider's revenue from a worker, essentially the withdrawal fees it collects, is of the same order as π^*_{welf} in the calibration, and the resource cost of the service is well below it, since EWA advances wages already earned and bears no credit risk. An employer buying the benefit for its workforce would moreover expect to pay less in aggregate than the sum of these fees, as programs are priced at the firm level, which is why the product can be offered to workers at no access fee.

Patience Panel B displays a deeper determinant of this slack. A patient worker facing the benchmark shock process self-insures: at $\delta = 0.999$ per week (an annualized 0.95),[8] she accumulates a buffer of about 0.9 paychecks, withdraws in only 8% of pay cycles, and values access at only 0.02% of the wage per cycle. As impatience rises, buffer-stock saving is crowded out and the product takes its place: at $\delta = 0.99$ per week the value of access reaches 0.28% of the wage per cycle, more than tenfold the patient level. EWA is thus a product for populations for whom holding liquid savings is costly or behaviorally difficult, consistent with its adoption patterns.

Shock exposure The value of access rises steeply with the size of mid-cycle expenditure shocks (Panel C): for the time-consistent type it nearly triples as the shock grows from 10% to 30% of the semimonthly wage. Insurance against larger shocks is exactly what self-insurance struggles to provide at reasonable buffer levels.

Present bias Present bias affects the value of access through two opposing channels. The *exposure channel* raises it: a present-biased worker holds a thinner buffer (0.62 versus 0.67 paychecks at the benchmark), reaches hand-to-mouth states more often, and therefore has more frequent recourse to the option, both against shocks and to align consumption with pay. The *dissipation channel* lowers it: her heavier, fee-bearing usage and the associated front-loading of consumption erode the long-run value of the option. Which channel dominates depends on how much there is to insure. At the benchmark shock size the two channels roughly offset, leaving the present-biased worker with a value close to that of the time-consistent one (0.097% against 0.109% of the wage per cycle); at larger shocks the exposure channel dominates and the present-biased worker values access *more* (Panel C, $\xi = 0.30$); at low risk aversion the dissipation channel dominates and the value of access falls to essentially zero (Panel E, $\gamma = 2$). Panel D adds the participation margin: the willingness to pay tracks the welfare value closely across the range of bias we consider, with a modest positive gap at the calibrated bias, and both peak at an intermediate bias ($\beta \approx 0.93$), where the thinner buffer maximizes the worker's exposure before fee dissipation takes over. The internality gap thus remains small throughout.

[8] Annualized discount factors compound over the model's year of 24 semimonthly pay cycles, that is, 48 model weeks.

Risk aversion The value of access increases with risk aversion for both types (Panel E), as expected of an insurance product, and more steeply for the present-biased type: curvature simultaneously disciplines her overuse and raises the value of insuring the exposure her thin buffer creates.

Taken together, the model predicts that EWA is most valuable for impatient, illiquid, risk-averse workers facing sizable expenditure shocks, which corresponds to the profile of the low-wage service workforces to which the product is marketed, and, among them, for moderately present-biased workers, whose thin buffers make the option most valuable whenever shocks are the relevant risk. As an illustration, a high-need worker combining impatience ($\delta = 0.99$) with large shocks ($\xi = 0.30$) values access at close to 0.4% of the wage per pay cycle, or about 10% of a semimonthly paycheck per year, an order of magnitude above the patient, well-buffered benchmark. This heterogeneity predicts that the benefits of EWA should be concentrated among lower-rank workers, and directly informs the targeting of the product.

[INSERT FIGURE 7 HERE]

5.5 Calibrating the Model to the Minu Usage Data

An advantage of our setting is that we observe a comprehensive history of the product usage over our sample period, which allows us to discipline the model's key parameters with the moments documented in Section 3 rather than assert them. Table 3, Panel A, summarizes the calibration.

The institutional parameters are set directly: the fee is $\eta = 0.98$, matching the roughly 2% per-withdrawal cost charged in practice, and the cap is $\bar{x} = 0.30$, the contractual 50% of earned wages evaluated at the withdrawal timing we observe. The minimum withdrawal is $x_{\min} = 0.015$, about one hundred pesos. Risk aversion is set to $\gamma = 3$, within the standard range of the household finance literature.

The remaining parameters, namely the weekly discount factor δ, the shock probability p and size ξ, and the present-bias parameter β of the frequent-user type, are set to match usage. The discount factor is identified by the crowding-out pattern of Panel B of Figure 7: patient workers self-insure and barely use the product, so matching the observed usage frequency of about 48% of user–pay cycles, and the thin liquid balances documented for

this population, requires an impatience wedge relative to the zero return on liquid savings; $\delta = 0.995$ per week (0.79 annualized) delivers both. Such impatience is the standard requirement for generating hand-to-mouth behavior in buffer-stock models. The implied annualized discount factor of 0.79 is within the range of field estimates of individual discount rates, which commonly exceed 25% per year (Warner and Pleeter, 2001, Harrison et al., 2002, Cohen et al., 2020), particularly among low-income and low-liquidity populations. Note also that δ and the present-bias parameter are separately identified: δ is pinned down by the behavior of the *average* user, whose usage frequency and thin buffer cannot be matched at higher values, while β is identified from the additional usage of the frequent-user group. Assuming a higher δ would therefore not weaken the role of present bias; it would make the average user's behavior unmatchable for a time-consistent type and shift more of the explanatory weight onto β. Given δ, we set the shock probability to $p = 0.30$, one shock every three or four cycles, and identify the shock size from the share of withdrawals that reach the cap. The share of withdrawals at the cap is monotone in ξ in the time-consistent benchmark, rising from 0.1% of withdrawals at $\xi = 0.10$ to 7.0% at $\xi = 0.30$, and it crosses the 0.73% observed in our data at $\xi = 0.13$, which we adopt. The at-cap moment is estimated with a bootstrap standard error of 0.12 percentage points, so this mapping locates ξ within a range rather than at a point, and the shock probability p is set rather than estimated; Section 5.4 reports how the value of access moves with both. A shock of that size is about 870 pesos at the mean semimonthly wage in our calibration sample: large enough to matter for a worker holding well under a paycheck of liquidity, and of the order of an out-of-pocket medical bill or an urgent repair. Finally, $\beta = 0.95$ is chosen so that the present-biased type reproduces the behavior of the frequent-user group, which uses the product in most pay cycles.[9]

Panel B of Table 3 compares model and data. The time-consistent type withdraws in 40% of pay cycles (48% in the data), withdrawing on average 9.7% of the cycle wage conditional on use (12.4% in the data), and reaching the cap in 0.7% of withdrawals against

[9]Regular usage alone does not identify present bias: a sufficiently impatient time-consistent worker (a lower δ) would also live hand-to-mouth and use the product in most cycles. The interpretation of frequent users as present-biased rests on the untargeted moments that the calibrated type reproduces, namely larger and more often capped withdrawals and a thinner buffer, and on the evidence of Section 4 that frequent users include higher-rank workers who are less likely to be financially constrained. We do not claim that this evidence rules out a sufficiently impatient time-consistent type, which would generate similar usage; what the framework requires is that some workers behave as if hand-to-mouth mid-cycle, and present bias is one economically natural way to reach that state.

0.7% in the data. The present-biased type behaves like the frequent users of Section 3: it uses EWA in 80% of cycles, withdraws more (12.3% of the wage, consistent with the larger and growing withdrawals of frequent users), reaches the cap in about 3% of withdrawals, and holds a substantially thinner buffer. The time-consistent type thus undershoots the intensive-margin moments of the average user, while the present-biased type overshoots the frequency. We read this as indicating a positive share of present-biased workers in the population, such that average usage intensity lies between the two columns. An untargeted moment provides validation: the model implies that 56% of the time-consistent type's usage (37% for the present-biased type) is shock-driven, the remainder occurring at hand-to-mouth states without a shock. Non-shock usage by the time-consistent type is not at odds with Proposition 1: it occurs precisely at states where the carryover constraint binds, which the time-consistent type reaches after histories of shocks have depleted her buffer. These shares bracket the survey evidence that unexpected and expected expenses account for roughly equal shares of usage purposes. The comparison is indicative rather than exact: the survey covers a later period and a broader set of firms than the administrative data, and reports purposes chosen by respondents rather than withdrawals driven by a shock. The concentration of withdrawals in the second week of the pay cycle, documented in Figure 1, is in contrast matched by construction, as the model makes withdrawals available only once wages have been earned mid-cycle.

[INSERT TABLE 3 HERE]

5.6 The Value of EWA for the Average Minu User

Panel C of Table 3 reports the pricing results. Three findings stand out.

First, EWA generates an economically meaningful welfare gain. The welfare-consistent access price is $\pi^*_{\text{welf}} = 0.11\%$ of the semimonthly wage per pay cycle for the time-consistent type and 0.10% for the present-biased type. This gain amounts to roughly 2.5% of one paycheck per year, or in capitalized terms to $\tau(\bar{m}) = 11.3\%$ of a paycheck for the time-consistent type (9.7% for the present-biased type) as a one-time equivalent at average liquidity. These values are computed under equilibrium behavior and are therefore already net of the withdrawal fees workers pay (0.08% and 0.20% of the wage per cycle for the two types): the time-consistent type's gain exceeds the fees she pays by about two fifths, and

the present-biased type, although her heavier usage generates a fee bill of about twice her gain, still retains her full surplus after fees. The gain also comes at close to zero resource cost, since withdrawal fees approximately cover the provider's cost and the product bears no credit risk. This value has two sources. The first is the direct smoothing of consumption around shocks shown in Figure 6, where mid-cycle consumption drops by 13% for the time-consistent type and 18% for the present-biased type without EWA, but stays nearly flat with it. The second, indirect, source operates through precautionary behavior: with EWA as a backstop, equilibrium payday buffers fall by 21% for the time-consistent type (from 0.67 to 0.53 paychecks) and by 30% for the present-biased type (from 0.62 to 0.43), freeing resources for consumption in normal times. This endogenous decline in liquid buffers is a testable prediction of the framework.

The framework also allows us to price the product's two services separately. Solving the model with the expenditure shock shut down ($p = 0$) isolates the pure consumption-timing service, whose stand-alone value is very limited. The time-consistent type essentially never uses EWA in that economy: absent shocks she is never caught liquidity-constrained, and the fee makes pure retiming unattractive, so the value of access is exactly zero. The naive present-biased type does use the product, in every cycle, withdrawing 3.7% of the wage to front-load consumption; but this usage has minimal effect on her long-run welfare and even her own payday self would decline free access ($\pi^*_{\text{WTP}} = 0$), since access only lets her mid-cycle self undo her plan at a fee.[10] The entire welfare gain reported above therefore derives from the interaction of the product with expenditure shocks: the timing service has no stand-alone value and becomes valuable only when shocks deplete liquidity. This decomposition also sharpens the reading of the survey evidence: although expected expenses account for roughly half of reported purposes of use, both the demand for EWA and its welfare value are rooted in its insurance function.

Second, the calibrated value should be read as the value for the *average* user in our setting, and related to the comparative statics of Section 5.4. The usage moments of the average Minu user imply moderate shocks and sufficient patience to sustain a buffer of about 0.7 paychecks, parameters at which the product operates well below the region where its

[10] At exactly $p = 0$, the deterministic intra-personal game admits a second, no-use stationary point for the present-biased type, an instance of the equilibrium indeterminacy arising under quasi-geometric discounting (Krusell and Smith, 2003). Because the usage equilibrium is the unique outcome at every $p > 0$ we examine, we select it as the $p \to 0$ limit; see Appendix A.

value peaks, though already more than five times the value it would have for a patient, well-buffered worker. For workers whose circumstances are more strained than the sample average, the same product is worth a multiple of the calibrated figure: more than 20% of a paycheck in capitalized terms below half-a-paycheck of liquidity (Panel A of Figure 7), nearly double the calibrated value when shocks reach 30% of the wage, the region in which present-biased workers gain the most, and close to 0.4% of the wage per cycle for the high-need profile of Section 5.4. The calibrated value is thus best interpreted as a lower bound on the product's potential in populations with less financial slack than our sample average, such as workforces facing more volatile expenses or weaker safety nets. Appendix Figure A.8 displays the underlying withdrawal policies: withdrawals provide partial insurance, increase as mid-cycle liquidity deteriorates, and remain below the cap except at severely depleted states; the present-biased type in addition withdraws at hand-to-mouth states without a shock, per Proposition 1.

Third, EWA does not monetize present bias. At the calibrated $\beta = 0.95$, the willingness to pay of the present-biased type exceeds its long-run welfare value by 0.015% of the semimonthly wage per cycle, about a sixth of that value: enrollment is not merely aligned with long-run welfare in sign, since both prices are positive and of the same order, but the wedge is small rather than absent. Panel D of Figure 7 traces both prices as bias deepens: the willingness to pay rises slowly and monotonically, from 0.110% at $\beta = 1$ to 0.120% at $\beta = 0.85$, while the welfare value falls, so that the wedge widens from nothing above $\beta = 0.975$ to two fifths of the welfare value at $\beta = 0.90$, and further at deeper bias. Deeper bias is not innocuous; rather, its potential cost is bounded by the product's design: because every advance is settled at the next payday and withdrawals are capped, the maximal damage from overuse is the fee on capped withdrawals, at most the fee on a capped withdrawal in every cycle, 0.6% of the wage. This bounded worst case is the key distinction between EWA and credit.

5.7 Comparison with Credit Access

To benchmark the welfare value of EWA against borrowing alternatives, we replace EWA with uncollateralized credit, keeping the income, shock, and preference structure identical. Credit differs from EWA in one structural dimension: it imposes no automatic repayment, so balances can roll over across pay cycles and accumulate interest. We consider two credit

regimes. The first, *payday-style credit*, reflects the short-term credit actually available to our population: a borrowing rate $R^{bor} = 1.05$ per week, in line with typical payday pricing of about 15% per two-week loan in the US and with the rates charged by Mexican pawnshops and informal lenders, a minimum loan size $B_{\min} = 0.10w$ (about US$35 at the mean wage) reflecting fixed origination costs, and a borrowing limit of $0.75w$. A rate far above the cost of EWA is not an assumption stacked against credit: an uncollateralized loan carries default risk that the lender must price, whereas EWA advances only wages already earned and is repaid at source, so it bears essentially no credit risk. The second, *frictionless cheap credit*, is a deliberately favorable benchmark: the rate $R^{bor} = 1.02$ per week matches EWA's effective cost and the minimum-size friction is removed, so that the only remaining difference with EWA is the absence of automatic repayment (and the looser limit). Comparing EWA with this benchmark isolates the pure effect of the repayment structure. Because the model is stationary, this comparison involves no end-of-horizon borrowing motive: all three regimes are evaluated in identical, cycle-invariant environments.

Table 4 and Figures 8 and 9 report the results. Payday-style credit does provide partial insurance in equilibrium, but far less than EWA and at a much higher price per unit of liquidity. The time-consistent type borrows in 4.7% of cycles and the present-biased type in 10.5%, taking loans just above the minimum size against shocks, and neither carries debt past payday. Access is worth 0.011% of the wage per cycle to both types, about a tenth of the value of EWA. Each unit of liquidity is also two and a half times dearer: measured against the average outstanding balance, interest runs at 5% per cycle where EWA's fee is 2%. Total interest paid, 0.027% and 0.065% of the wage per cycle, is nonetheless smaller than the fees paid under EWA, precisely because so little borrowing takes place. The comparison is instructive on both margins: the rate makes each unit of liquidity dearer, and the minimum loan size forces workers to buy more liquidity than they need, so that credit is used only against larger shortfalls. EWA, which is divisible down to the peso and cheaper per unit, dominates this functioning market alternative, consistent with the survey evidence that workers substitute away from these alternatives upon EWA adoption.

The frictionless benchmark isolates the role of the repayment structure. For the time-consistent type, cheap flexible credit and EWA are worth the same amount, 0.109% of the wage per cycle: flexibility buys nothing when self-control is not an issue, and the automatic repayment of EWA sacrifices nothing. EWA thus provides time-consistent workers with

the value of a frictionlessly priced credit line, although it requires no underwriting or credit registry and bears no default risk. The same holds for a *sophisticated* present-biased worker, who correctly anticipates her future behavior: facing her paycheck and compounding interest at payday, her own bias-weighted tradeoff makes repayment optimal, so she carries almost no debt across paydays, and the automatic repayment of EWA merely enforces what she would do anyway. The two products are worth 0.103% and 0.106% to her, a gap we do not read as a ranking.[11]

The economically decisive case is instead the one in which workers are *naive* about their own bias, that is, when they believe each period that their future selves will behave in a time-consistent manner. The evidence on repeated payday-loan rollovers supports this assumption for a sizable fraction of short-term borrowers (Melzer, 2011). A naive worker borrows on flexible credit expecting her payday self to repay; her payday self, subject to the same bias she did not anticipate, rolls part of the balance over. Panel B of Figure 9 traces the consequences: at the calibrated bias the naive worker still repays nearly everything (rollover of 0.009% of the wage) and the two products remain close, around 0.10% each, but as bias deepens they diverge sharply. At $\beta = 0.85$, EWA is worth 0.036% of the wage per cycle while frictionless credit already reduces welfare (-0.019%), with rollover balances of 3.0% of the wage; at $\beta = 0.80$, the cost of credit access reaches 0.16% per cycle as rollover reaches 11% of the wage and interest payments compound on balances that each payday self never planned to carry, a debt trap in which the household is inefficiently stuck rolling over interest-bearing debt, while EWA is still worth a positive 0.027%; at $\beta = 0.75$, credit costs 0.28% of the wage per cycle against 0.11% for EWA, which by that depth of bias is itself costly, but only a third as costly. The asymmetry follows from the contract structure: EWA *disciplines* the borrowing of naive present-biased workers through its forced repayment. Under EWA, the automatic paycheck deduction executes the repayment the naive worker wrongly assumes she will make, and the cap limits her exposure, so the worst case is the fee on capped withdrawals. Under flexible credit, by contrast, the gap between

[11] The gap is of the same order as the numerical precision with which the two products can be compared: an access price is a difference of value functions of order 10^{-3}, and the EWA and credit columns are computed by methods that price EWA within 5% of each other (see the notes to Table 4). Separately, the sophisticated present-biased equilibrium is an intra-personal fixed point, and it is the one case in which the solution need not be unique: at the calibrated $\beta = 0.95$ it converges, but at deeper bias the value selected depends on the starting point. That is why the bias sweep in Panel B of Figure 9 is drawn for the naive type, whose problem involves no such fixed point—she plans against a time-consistent continuation and her behavior is a single best response to it—and is numerically stable at any depth of bias.

planned and actual repayment is left to the worker's biased payday self and compounds at interest. EWA thus weakly dominates even frictionlessly priced flexible credit for present-biased workers who underestimate their bias, and the dominance grows with the bias. This disciplining role is structural rather than behavioral: it operates through the contract, not through any change in the worker's awareness of her bias. The repayment structure is therefore a central design feature of EWA rather than a limitation: it makes the product's value robust to the very bias that generates the heaviest demand for it, whether or not workers recognize that bias in themselves.[12]

This asymmetry is clearest once we separate what a worker expects to gain from credit access, ex ante, from what she realizes ex post. The *ex-ante* willingness to pay for flexible credit is nearly the same for naive and sophisticated present-biased workers (0.113% of the wage per cycle at the calibrated bias for the naive worker, 0.106% for the sophisticate): each evaluates access under the belief that her payday self will repay, and so values the flexibility almost identically. Their *ex-post* realized welfare, however, differs, and this is precisely where EWA's forced repayment does its work. Define the ex-post gain from EWA over credit as the difference in realized long-run welfare between the two products under each type's actual behavior. For the sophisticated present-biased worker this gain is essentially zero at the calibrated bias, because her realized repayment matches her plan and the two products deliver the same path. For the naive worker it is positive and rises steeply with the depth of bias: about 0.055% of the wage per cycle at $\beta = 0.85$ and 0.185% at $\beta = 0.80$, as her realized borrowing under credit accumulates the rollover debt she did not anticipate while EWA mechanically prevents it. Two workers with the same willingness to pay for credit therefore obtain very different realized value from being given EWA instead: the naive worker, whose realized behavior departs most from her plan, gains the most.

[INSERT FIGURES 8 and 9 HERE]

[INSERT TABLE 4 HERE]

[12]Naivete also widens the internality gap, since willingness to pay reflects usage plans that are not followed: at the calibrated bias, the naive gap is 0.015 percentage points for EWA and 0.018 for flexible credit, against 0.001 for the sophisticated present-biased worker. Even under naivete, therefore, EWA's appeal remains closely aligned with its long-run value. Naive usage of EWA is broadly similar to sophisticated usage (withdrawals in 79.7% of cycles versus 72.8%, averaging 12.3% and 12.8% of the wage respectively), so our calibration is robust to the sophistication assumption.

5.8 Comparison with More Frequent Pay

A natural alternative to EWA is to pay workers more often, since a shorter pay cycle relieves the same income–expenditure mismatch directly. We solve the model under weekly pay, splitting the semimonthly wage into two equal payments. More frequent pay and EWA are partial substitutes: under weekly pay the welfare value of EWA falls from 0.109% to 0.043% of the wage per cycle for the time-consistent type and its usage falls from 40% to 11% of cycles, with the present-biased type behaving similarly. Weekly pay is also worth more on its own, 0.46% of the wage per cycle, because it advances half a paycheck every cycle rather than a capped fraction drawn only when needed; this reflects the scale of the intervention, not an inefficiency of EWA. It also creates no rollover debt, so it is worth as much to naive present-biased workers (0.46%) as to sophisticated ones (0.48%) and does the naive worker no harm. The automatic repayment that separates EWA from credit is therefore not what separates it from more frequent pay: both are safe for the present-biased.

The behavior of weekly-paid workers in our sample provides a partial test of the substitution effect predicted by the model. Although demand for Early Wage Access (EWA) persists under a weekly pay schedule—a noteworthy finding—it is quantitatively lower than that observed among semimonthly-paid workers, with weekly-paid workers utilizing EWA in 31.8% of pay cycles compared to 57.8% for semimonthly-paid workers. This pattern aligns with the model's directional prediction, albeit less sharply than the predicted difference of 40% versus 11%. It is important to note that pay frequency is not randomly assigned; weekly-paid workers are generally poorer and face greater financial constraints, factors that counteract the substitution effect isolated by the experiment. Moreover, their withdrawal patterns are largely dictated by shorter pay cycles and contractual limitations rather than heightened liquidity demand. Specifically, the calibrated shock size corresponds to approximately one-tenth of a semimonthly paycheck but constitutes a substantial portion of a weekly paycheck. Additionally, the 50% cap on wages earned but not yet paid is correspondingly lower for weekly-paid workers, often binding and restricting early withdrawals. For example, a worker requiring 30% of her weekly paycheck cannot access these funds until around the fifth day of the week. Consequently, the timing of withdrawals is strongly influenced by contractual pay schedule constraints rather than immediate liquidity needs. These considerations justify calibrating the model primarily on semimonthly-paid

workers.

Why, then, might an employer prefer EWA? On the worker's side the model gives at least as much to more frequent pay, so the case for EWA rests on the firm's side, which the model does not price. Paying more often multiplies the fixed costs of payroll, advances the entire wage bill rather than the fraction workers draw, and is often constrained by law, collective agreements, or legacy payroll systems. EWA layers onto any existing schedule at close to zero incremental cost and, being state-contingent, moves far less cash while capturing much of the benefit. It is thus a low-cost, on-demand approximation to higher pay frequency, complementing work on the optimal frequency of pay (Parsons and Van Wesep, 2013, Baugh and Correia, 2022).

5.9 Implications for EWA Product Design

The model's results have several implications for the design of EWA programs. The contractual cap at 50% of earned wages accommodates essentially all insurance-motivated withdrawals, as equilibrium withdrawals reach the cap in only about 1% to 3% of uses, while bounding the exposure of present-biased users. The withdrawal fee disciplines usage beyond cost recovery: by Proposition 1, any strictly positive fee eliminates non-shock usage by workers with interior buffers, so a minimal fee preserves the full insurance value while deterring gratuitous front-loading; fee levels far above cost, in contrast, would erode the insurance value that justifies the product. The withdrawal fee is also a margin for the employer rather than a fixed feature of the product, since the firm can absorb it on the worker's behalf, which in the model corresponds to setting $\eta = 1$.[13] The automatic repayment structure bounds the harm from overuse and differentiates EWA from credit for the present-biased users who use the product most intensively, all the more so when those users underestimate their own bias, in which case flexible credit reduces welfare while the worst case under EWA remains close to zero. Finally, the concentration of the product's

[13]The erosion is quantitatively steep: recomputing the calibrated model at a 4% proportional fee, usage falls materially for both types and the welfare value of access falls to 0.07% and 0.04% of the wage per cycle, roughly two thirds and two fifths of its benchmark level. Conversely, when the employer absorbs most of the fee, at $\eta = 0.995$, the value of access rises to 0.22% of the wage per cycle for the time-consistent type and 0.21% for the present-biased type, about twice their benchmark levels. The internality gap does not widen as the fee falls: it stays at 0.014% of the wage, so it shrinks from a sixth of the welfare value at the benchmark fee to a fifteenth. What limits the exposure of present-biased workers is therefore the repayment structure and the cap rather than the fee, which is why a program in which the employer bears the fee does not invite the overuse the fee might be thought to deter.

value at low liquidity implies that EWA delivers most where workforces are low-wage and low-buffer. That is a statement about the value of access rather than about who takes it up, and the two differ in our data, as adoption rises with job rank while usage intensity falls with it.

The model's precautionary-savings prediction also carries a cautionary implication for the regulatory debate. Access endogenously lowers liquid buffers, by roughly 20% to 30% at the calibration, so the observation that EWA users hold less savings than non-users should not be read as evidence that the product undermines financial health: the model generates exactly this pattern as the equilibrium response to being insured, and the decline in self-insurance is part of the welfare gain rather than a symptom of distress. By the same logic, the thin balances observed among users in the cross-section partly reflect the product itself, and not only the selection of illiquid workers into it.

Taken together, these results give a bounded answer to the concern that EWA carries unintended costs. Three potential costs are present but limited. First, the product is not cheap, as its fees imply high annualized rates; yet for the average user in our calibration the welfare gain exceeds the fees paid, and the concentration of value at low liquidity means that the workers who use it most are those for whom it is worth most. Second, EWA could in principle monetize present bias. The model implies that it does so only to a limited extent at moderate degrees of bias, where willingness to pay tracks long-run welfare closely and the internality gap is a small fraction of the value of access, although that fraction grows as bias deepens. Third, heavy use by present-biased workers could erode welfare, yet the automatic repayment and the cap bound the worst case at the fee on a capped withdrawal in every cycle, 0.6% of the wage, unlike the unbounded losses credit can inflict on the same workers. These conclusions are model-based: they establish that the costs of EWA are structurally limited rather than measuring them directly, and they leave open questions our data cannot address, such as whether availability fosters dependence over longer horizons, whether reduced turnover forecloses efficient job moves, and how the product is valued by the many workers who never adopt it. We view credible empirical evidence on these margins as an important avenue for future research.

6 Earned Wage Access and Employee Retention

The framework of Section 5 quantifies economically meaningful welfare gains from EWA access, concentrated among liquidity-constrained workers. This section takes these predictions to the data. If access generates the surplus the model implies, workers should value employers that provide it, and this attachment should translate into lower employee turnover, especially while the share of firms offering the benefit remains low, and disproportionately so among the workers for whom the model implies the largest gains. Employee retention thus provides a market-based check on the welfare significance of EWA, and identifies a source of economic gain for employers, as turnover is costly and unusually high in low-skill service industries. We emphasize at the outset that the analysis that follows is associational: we ask whether the retention patterns in the data are consistent with what the framework implies, not whether EWA causes workers to stay.

6.1 Main Result

We begin with two descriptive comparisons, reported in Appendix Figures A.9 and A.10. The distribution of completed tenure among leavers is shifted toward longer spells for EWA users, and among workers hired shortly before or after the introduction of EWA, users remain employed at higher rates, by about five percentage points eight months out. We read these comparisons only as suggestive, because they are vulnerable to a mechanical bias: a worker who leaves soon after joining has little opportunity to adopt EWA and is therefore classified as a non-user precisely *because* she left quickly, which by construction associates non-use with shorter tenure.[14] We therefore rely on the panel hazard model and the propensity-based design that follow, rather than on these unconditional comparisons, for our substantive results.

This raw data correlation could however be driven by selection effects, for instance because characteristics that predict EWA adoption are correlated with turnover rate. To mitigate this concern, we estimate the relationship between using earned wage access and the likelihood of leaving the firm while controlling for observable characteristics of workers.

[14]This is a form of immortal-time bias. The hazard specification below addresses it by treating EWA use as a time-varying status that switches on only from the pay cycle of first use, so that pre-adoption worker-cycles are not counted as user time; the propensity-based design further avoids conditioning on realized use altogether.

We use a Cox proportional hazard model, with the following specification:

$$\lambda(t \mid X_i, Used\ EWA_{i,t}) = \lambda_{0,s(i)}(t) \times \exp\left(\beta_1 Used\ EWA_{i,t} + \Gamma' X_i\right)$$

where $\lambda(t \mid \cdot)$ is the hazard of worker i leaving the firm in half-month t, with t measured as half-months since the worker's company introduced EWA (workers hired after the introduction enter the risk set at hiring), $Used\ EWA_{i,t}$ is an indicator equal to one from the half-month in which the worker first uses EWA onward, and zero before, and X_i collects time-invariant worker characteristics (gender in column 3; gender and age in column 4). The baseline hazard $\lambda_{0,s(i)}(t)$ is stratified by the cell $s(i)$ to which worker i belongs: company $\times$ pay frequency in column 1, additionally the worker's entry quarter in column 2, and additionally the worker's job rank in columns 3 and 4. Stratification leaves the baseline hazard of each cell unrestricted, so pay-frequency and cohort differences in turnover are absorbed. Standard errors are clustered at the employee level.

Table 5 presents the estimated coefficients, which are log hazard ratios: a coefficient $\hat{\beta}$ implies a proportional change of $e^{\hat{\beta}} - 1$ in the separation hazard. In columns 1 and 2, stratifying on company and pay frequency and then also on starting-quarter cohorts, EWA users on average have about a 10 to 11 percent lower probability of leaving the company in a given pay cycle than non-users. When we additionally stratify on job rank and control for gender in column 3, and add age in column 4, the association strengthens slightly and remains economically and statistically significant: EWA users exhibit a 12 percent lower separation hazard with the full set of controls.

[INSERT TABLE 5 HERE]

Next, we run split-sample regressions by gender, job rank, and age group to uncover potential heterogeneity in the relationship between EWA usage and worker retention in Table 6. First, the association is statistically significant among women, for whom being an EWA user is associated with an 18 percent lower separation hazard, and is smaller and insignificant among men; the two point estimates are not, however, formally compared, and the framework offers no prediction along this dimension. Second, splitting workers along their skill level in columns 3 and 4 delivers estimates of very similar size for the two groups, a 13 percent lower hazard for low-rank workers and a 13 percent lower hazard

for mid- and high-rank workers, both statistically significant. The retention correlation is thus not concentrated in either rank group. We do not read this cross-sectional split as informative about the framework in either direction: rank is a coarse proxy for liquidity, and the frequent users whose behavior the model reads as present-biased include higher-rank workers, for whom the framework also implies meaningful gains.

The sharpest contrast is by age. Young users are associated with a 19 percent lower separation hazard, whereas for older workers the point estimate is zero. The two estimates are not formally compared, and age is not a dimension along which the framework makes a prediction, so we report this contrast as a feature of the data rather than as evidence on the model.

[INSERT TABLE 6 HERE]

To further address the concern that actual EWA adoption is endogenous to worker turnover propensity, we implement a propensity-based difference-in-differences (DiD) design. We first estimate each worker's predicted probability of EWA adoption using a logit model based solely on pre-determined worker characteristics (gender, age, job rank, and starting month). Workers are then classified as "High Propensity" if their predicted adoption probability exceeds the sample median. The assignment uses only characteristics determined before EWA is introduced, so it is unaffected by a worker's later usage decision, although it does not by itself deliver exogenous variation: the design compares workers who differ in observable characteristics, and identification requires that their separation rates would have moved in parallel absent the benefit. The DiD specification regresses an indicator for leaving the firm within a given horizon (2, 4, 6, and 8 months) on the interaction of a post-EWA-introduction indicator with the High Propensity group indicator, along with company fixed effects and company-by-post fixed effects to absorb time-varying firm-level shocks. Table 7 shows the interaction coefficient Post $\times$ High Propensity To Adopt is negative at every horizon, ranging from -4.3 percentage points at the 2-month horizon (significant at the 5% level) to -5.8 percentage points at the 8-month horizon (significant at the 10% level); the 6-month estimate is of similar magnitude but imprecisely estimated. These patterns are consistent with EWA access being associated with lower turnover among the workers most likely to use the product. These estimates are not directly comparable in magnitude with the Cox results above, since they compare groups defined by predicted

rather than actual adoption, and we read them only as a further pattern consistent with the framework rather than as a causal estimate. Standard errors are clustered at the start-quarter level to account for within-cohort correlation.

[INSERT TABLE 7 HERE]

6.2 Robustness

The half-month grid on which we measure separations carries heavy ties, which the Cox partial likelihood handles only approximately. Appendix Table A.1 therefore re-estimates the extensive-margin specification as a discrete-time hazard on the employee–half-month risk-set rows, using a complementary log-log link, the exact grouped-time analogue of the proportional-hazards model, and a logit link. Coefficients are reported in exponentiated form, so they are directly comparable to $e^{\hat{\beta}}$ from Table 5. The estimates line up with the Cox results in magnitude: with start-quarter dummies and worker characteristics, EWA use is associated with a separation hazard about 11% lower (0.894 in column 4), against the $e^{-0.128} = 0.88$ implied by column 4 of Table 5. Without cohort controls the association is larger, a hazard 22% lower, and it is statistically significant in every column. It is important to note that start quarter enters here as dummies rather than as strata, a coarser cohort control than in Table 5.

Finally, we ask whether the retention correlation is concentrated among the workers whose usage pattern the model associates with the tightest liquidity constraints. We classify each user over the first six pay cycles following her first withdrawal: a *Regular* user is active in at least four of those six cycles, and an *Other* user, the pooled remainder, withdraws only occasionally. Appendix Table A.2 re-estimates the hazard model on the full sample with time-varying group membership that switches on once the classification window closes, so that a user's pre-classification person-time contributes to the comparison group rather than to her eventual group. The Regular coefficient is negative in every column and, in column 3, which controls for gender, age, and monthly salary, corresponds to a 26 percent lower separation hazard, whereas occasional users are statistically indistinguishable from non-users throughout. The retention correlation therefore sits with the regular users, that is, with the group whose behavior the model reads as hand-to-mouth, which is the pattern the framework implies. We note that the estimate is significant in the specification that

controls for monthly salary and not in the two sparser ones, and read it accordingly.

Taken together, the hazard, discrete-time, and propensity-based results consistently associate the usage of earned wage access with meaningfully lower turnover, and this association is concentrated among the regular users whose behavior the framework reads as hand-to-mouth, which is where it implies the largest gains. These estimates remain associations: our designs mitigate selection concerns without eliminating them, and we do not interpret them as causal effects of EWA on turnover.

The framework also disciplines the interpretation of these magnitudes. The steady-state welfare gain from EWA access, of the order of 0.1% of the wage as a flow or a tenth of a paycheck capitalized, is meaningful but modest relative to the 10–13% lower separation probability we document. Our separation measure also pools quits and dismissals, which the HR records do not distinguish. If financial strain raises the likelihood of dismissal, as absenteeism and job performance respond to it, then workers in strained circumstances separate more often for reasons unrelated to their valuation of the benefit, which works against finding the association we document. Two readings are consistent with this gap. First, part of the retention correlation reflects selection, as more attached workers adopt the benefit; our designs mitigate but cannot eliminate this. Second, the model values EWA under equilibrium behavior, in which shocks are partially self-insured; the state-contingent value $\tau(m)$ exceeds 20% of a paycheck in the depleted-liquidity states in which a worker's attachment to the employer is most at risk, that is, when a financial emergency would otherwise trigger informal borrowing, moonlighting, or a move to the job offering the most immediate cash. A benefit whose value is concentrated in those states can plausibly move separation decisions by more than its unconditional flow value suggests, at close to zero cost to the employer. We view the joint quantification of this retention channel as a promising avenue for future research.

7 Conclusion

Using novel usage and survey data from a Mexican FinTech firm, we study the use by workers of earned wage access, an innovative service increasingly offered by employers as a benefit. We document that the product meets significant demand, and that its usage is associated with higher employee retention, a pattern consistent with the welfare gains

our framework implies although we cannot rule out selection into adoption. We develop a stationary buffer-stock framework, disciplined by our usage data, in which EWA provides liquidity insurance and consumption-timing alignment. The framework rationalizes the usage patterns we observe, quantifies welfare gains that are concentrated among low-liquidity workers and approach the value of a frictionlessly priced credit line, and shows that the automatic repayment structure of EWA makes these gains more robust to present bias than those from credit alternatives.

References

Baugh, Brian, and Filipe Correia, 2022, Does paycheck frequency matter? evidence from micro data, *Journal of Financial Economics* 143, 1026–1042.

Baugh, Brian, Jesse B Leary, and Jialan Wang, 2018, When is it hard to make ends meet?, Working paper, National Bureau of Economic Research.

Berg, Tobias, Valentin Burg, Jan Keil, and Manju Puri, 2025, The Economics of "Buy Now, Pay Later": A Merchant's Perspective, *Journal of Financial Economics* 171, 104093.

Berg, Tobias, Andreas Fuster, and Manju Puri, 2022, FinTech Lending, *Annual Review of Financial Economics* 14, 187–207.

Bloom, Nicholas, Ruobing Han, and James Liang, 2024, Hybrid working from home improves retention without damaging performance, *Nature* 630, 920–925.

Bode, Christiane, Jasjit Singh, and Michelle Rogan, 2015, Corporate Social Initiatives and Employee Retention, *Organization Science* 26, 1702–1720.

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman, 2015, The Value of Hiring through Employee Referrals*, *The Quarterly Journal of Economics* 130, 805–839.

Calvet, Laurent E., Claire Célérier, Paolo Sodini, and Boris Vallée, 2023, Can Security Design Foster Household Risk-Taking?, *The Journal of Finance* 78, 1917–1966.

Carnahan, Seth, David Kryscynski, and Daniel Olson, 2017, When Does Corporate Social Responsibility Reduce Employee Turnover? Evidence from Attorneys Before and After 9/11, *Academy of Management Journal* 60, 1932–1962.

Carvalho, Leandro S, Stephan Meier, and Stephanie W Wang, 2016, Poverty and economic decision-making: Evidence from changes in financial resources at payday, *American Economic Review* 106, 260–284.

Chava, Sudheer, Rohan Ganduri, Nikhil Paradkar, and Yafei Zhang, 2021, Impact of marketplace lending on consumers' future borrowing capacities and borrowing outcomes, *Journal of Financial Economics* 142, 1186–1208.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White, 2020, Measuring Time Preferences, *Journal of Economic Literature* 58, 299–347.

Consumer Financial Protection Bureau, 2024, Data Spotlight: Developments in the Paycheck Advance Market, Technical report, Consumer Financial Protection Bureau.

Demirgüç-Kunt, Asli, Leora Klapper, Dorothe Singer, and Saniya Ansar, 2022, *The Global Findex Database 2021: Financial Inclusion, Digital Payments, and Resilience in the Age of COVID-19* (World Bank, Washington, DC).

Di Maggio, Marco, and Vincent Yao, 2021, Fintech Borrowers: Lax Screening or Cream-Skimming?, *The Review of Financial Studies* 34, 4565–4618.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo, 2018, The Economic Consequences of Hospital Admissions, *American Economic Review* 108, 308–352.

Dube, Arindrajit, T. William Lester, and Michael Reich, 2016, Minimum Wage Shocks, Employment Flows, and Labor Market Frictions, *Journal of Labor Economics* 34, 663–704.

Emanuel, Natalia, and Emma Harrington, 2024, The Payoffs of Higher Pay: Labor Supply and Productivity Responses to a Voluntary Firm Minimum Wage, Staff report no. 1182, Federal Reserve Bank of New York.

Fink, Günther, B. Kelsey Jack, and Felix Masiye, 2020, Seasonal Liquidity, Rural Labor Markets, and Agricultural Production, *American Economic Review* 110, 3351–3392.

Gathergood, John, Benedict Guttman-Kenney, and Stefan Hunt, 2019, How Do Payday Loans Affect Borrowers? Evidence from the U.K. Market, *The Review of Financial Studies* 32, 496–523.

Ghosh, Pulak, Boris Vallée, and Yao Zeng, 2026, FinTech Lending and Cashless Payments, *The Journal of Finance* 81, 1053–1101.

Harrison, Glenn W., Morten I. Lau, and Melonie B. Williams, 2002, Estimating Individual Discount Rates in Denmark: A Field Experiment, *American Economic Review* 92, 1606–1617.

Haushofer, Johannes, and Ernst Fehr, 2014, On the psychology of poverty, *Science* 344, 862–867.

Indarte, Sasha, Raymond Kluender, Ulrike Malmendier, and Michael Stepner, 2026, Consumption Wedges: Measuring and Diagnosing Distortions, Working Paper 34891, National Bureau of Economic Research.

Iskhakov, Fedor, Thomas H. Jørgensen, John Rust, and Bertel Schjerning, 2017, The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks, *Quantitative Economics* 8, 317–365.

Kaur, Supreet, Sendhil Mullainathan, Suanna Oh, and Frank Schilbach, 2025, Do Financial Concerns Make Workers Less Productive?, *The Quarterly Journal of Economics* 140, 635–689.

Kluender, Raymond, Neale Mahoney, Francis Wong, and Wesley Yin, 2025, The Effects of Medical Debt Relief: Evidence from Two Randomized Experiments, *The Quarterly Journal of Economics* 140, 1187–1241.

Krusell, Per, and Anthony A. Smith, 2003, Consumption–savings decisions with quasi-geometric discounting, *Econometrica* 71, 365–375.

Laibson, D., 1997, Golden Eggs and Hyperbolic Discounting, *The Quarterly Journal of Economics* 112, 443–478.

Maxted, Peter, 2025, Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma, *The Quarterly Journal of Economics* 140.

Meier, Stephan, and Charles D. Sprenger, 2015, Temporal Stability of Time Preferences, *Review of Economics and Statistics* 97, 273–286.

Melzer, Brian T., 2011, The Real Costs of Credit Access: Evidence from the Payday Lending Market*, *The Quarterly Journal of Economics* 126, 517–555.

Morse, Adair, 2011, Payday lenders: Heroes or villains?, *Journal of Financial Economics* 102, 28–44.

O'Donoghue, Ted, and Matthew Rabin, 2015, Present Bias: Lessons Learned and To Be Learned, *American Economic Review* 105, 273–279.

Olafsson, Arna, and Michaela Pagel, 2018, The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software, *The Review of Financial Studies* 31, 4398–4446.

Oyer, Paul, 2004, Why Do Firms Use Incentives That Have No Incentive Effects?, *The Journal of Finance* 59, 1619–1650.

Oyer, Paul, and Scott Schaefer, 2005, Why do some firms give stock options to all employees?: An empirical examination of alternative theories, *Journal of Financial Economics* 76, 99–133.

Parker, Jonathan A, 1999, The Reaction of Household Consumption to Predictable Changes in Social Security Taxes, *American Economic Review* 89, 959–973.

Parsons, Christopher A., and Edward D. Van Wesep, 2013, The timing of pay, *Journal of Financial Economics* 109, 373–397.

Reher, Michael, and Stanislav Sokolinski, 2024, Robo advisors and access to wealth management, *Journal of Financial Economics* 155, 103829.

Rust, John, 1987, Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher, *Econometrica* 55, 999–1033.

Shapiro, Jesse M., 2005, Is there a daily discount rate? Evidence from the food stamp nutrition cycle, *Journal of Public Economics* 89, 303–325.

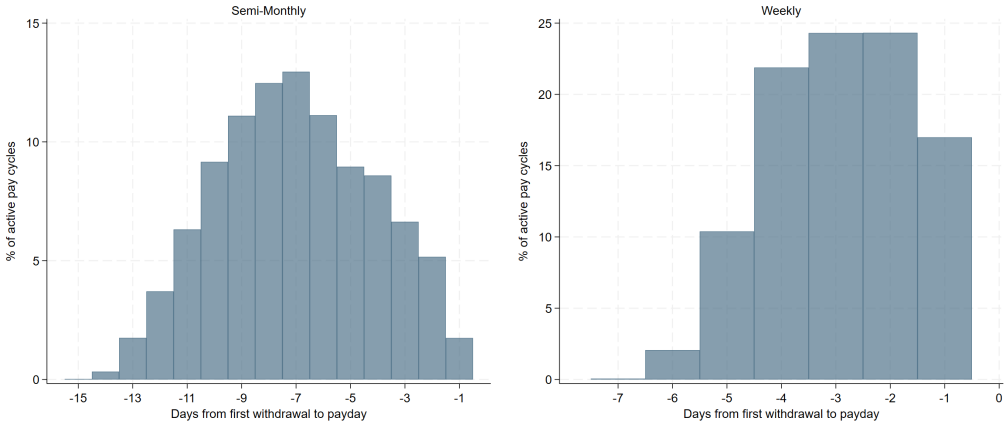
Stephens, Melvin, 2003, "3rd of tha Month": Do Social Security Recipients Smooth Consumption Between Checks?, *American Economic Review* 93, 406–422.

Visa, 2019, Earned Wage Access, *Visa Insights* .

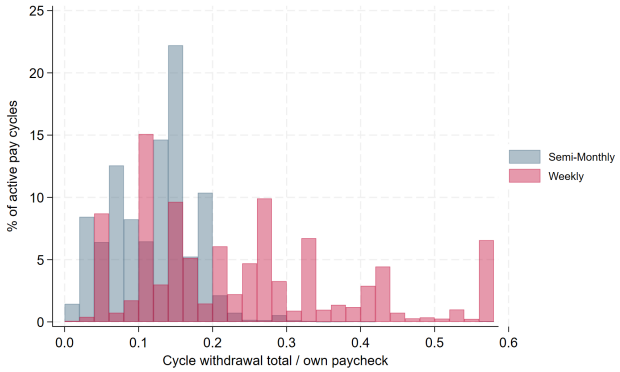
Warner, John T., and Saul Pleeter, 2001, The Personal Discount Rate: Evidence from Military Downsizing Programs, *American Economic Review* 91, 33–53.

Figures

Panel A: Withdrawal timing within the pay cycle



Panel B: Withdrawal size relative to the paycheck



Panel C: Withdrawal size relative to available income

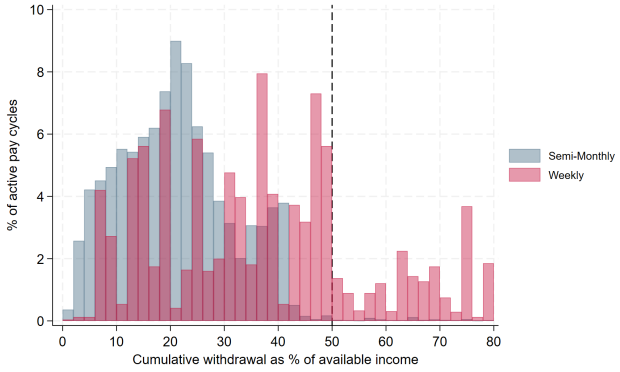


Figure 1: **Earned Wage Access: Withdrawal Timing and Size**

Notes: These figures present the timing and the scaled size of EWA withdrawals, separately for workers paid semimonthly and workers paid weekly. Panel A plots the distribution of the number of days between the worker's first withdrawal in a pay cycle and the next payday. The left sub-panel plots the semimonthly-paid workers and the right sub-panel plots the weekly paid workers. Panel B plots the distribution of the amount of withdrawals scaled by the worker's own paycheck, at pay-cycle level. Panel C plots the distribution of the cumulative amount withdrawn within a pay cycle as a percentage of the wages the worker had earned by the time of her last request in that cycle, at the worker–pay cycle level; the dashed vertical line marks the cap of 50% of earned wages. Panel C is winsorized at the 99th percentile.

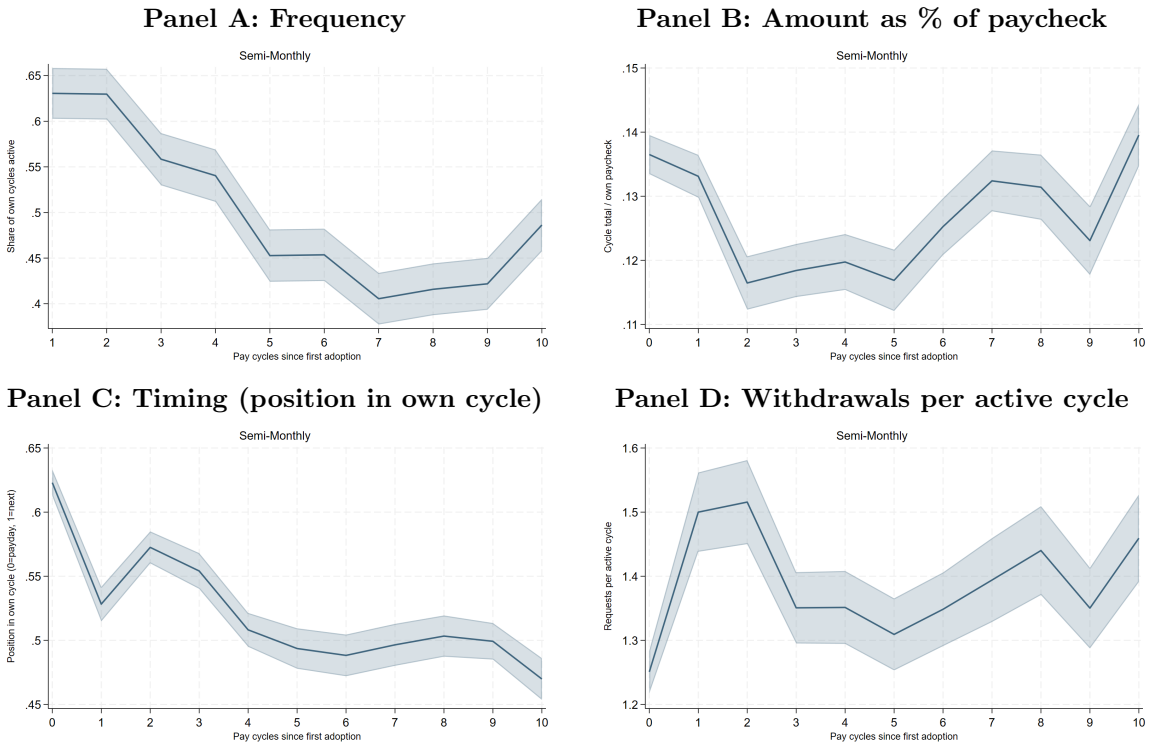


Figure 2: **Average Withdrawal Behaviors over Time**

Notes: These figures present the evolution of the frequency, amount, timing and intensity of withdrawals over the pay cycles following adoption, for EWA users paid semimonthly. The horizontal axis counts a worker's own pay cycles since her first withdrawal. The panel is balanced: it retains the users who are observable for at least ten pay cycles after adoption. Frequency is the share of a worker's own cycles in which she makes at least one withdrawal; the three remaining panels condition on cycles with at least one withdrawal and report, respectively, the cycle total scaled by the worker's own paycheck, the position of the first request within the cycle (0 = payday, 1 = the next payday), and the number of requests per active cycle. Shaded bands are 95% confidence intervals from standard errors clustered on the worker. The corresponding figures for weekly-paid workers are reported in Appendix Figure A.2.

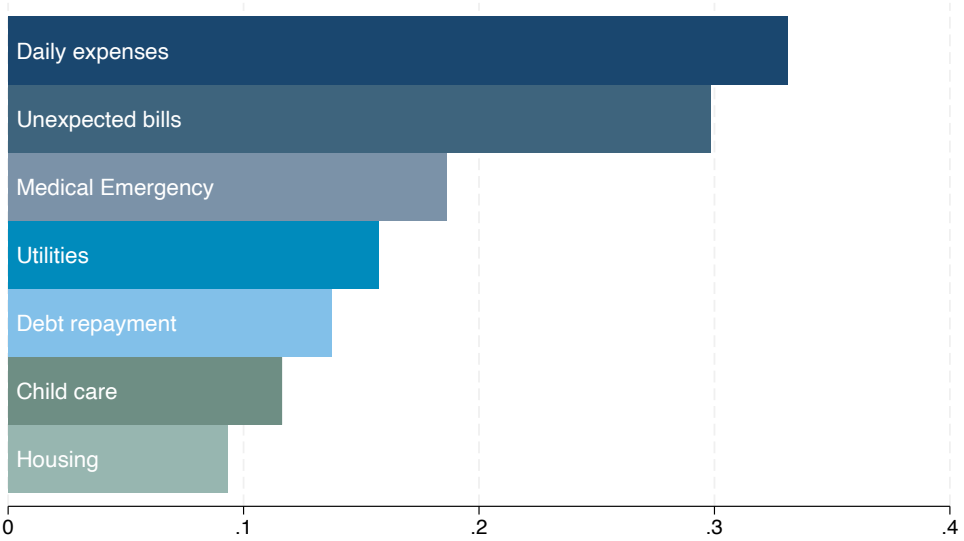


Figure 3: **EWA Purposes of Use**

Notes: This figure plots the proportions of top three answers chosen by respondents to survey question: "What are the main purposes that you use Minu EWA for?".

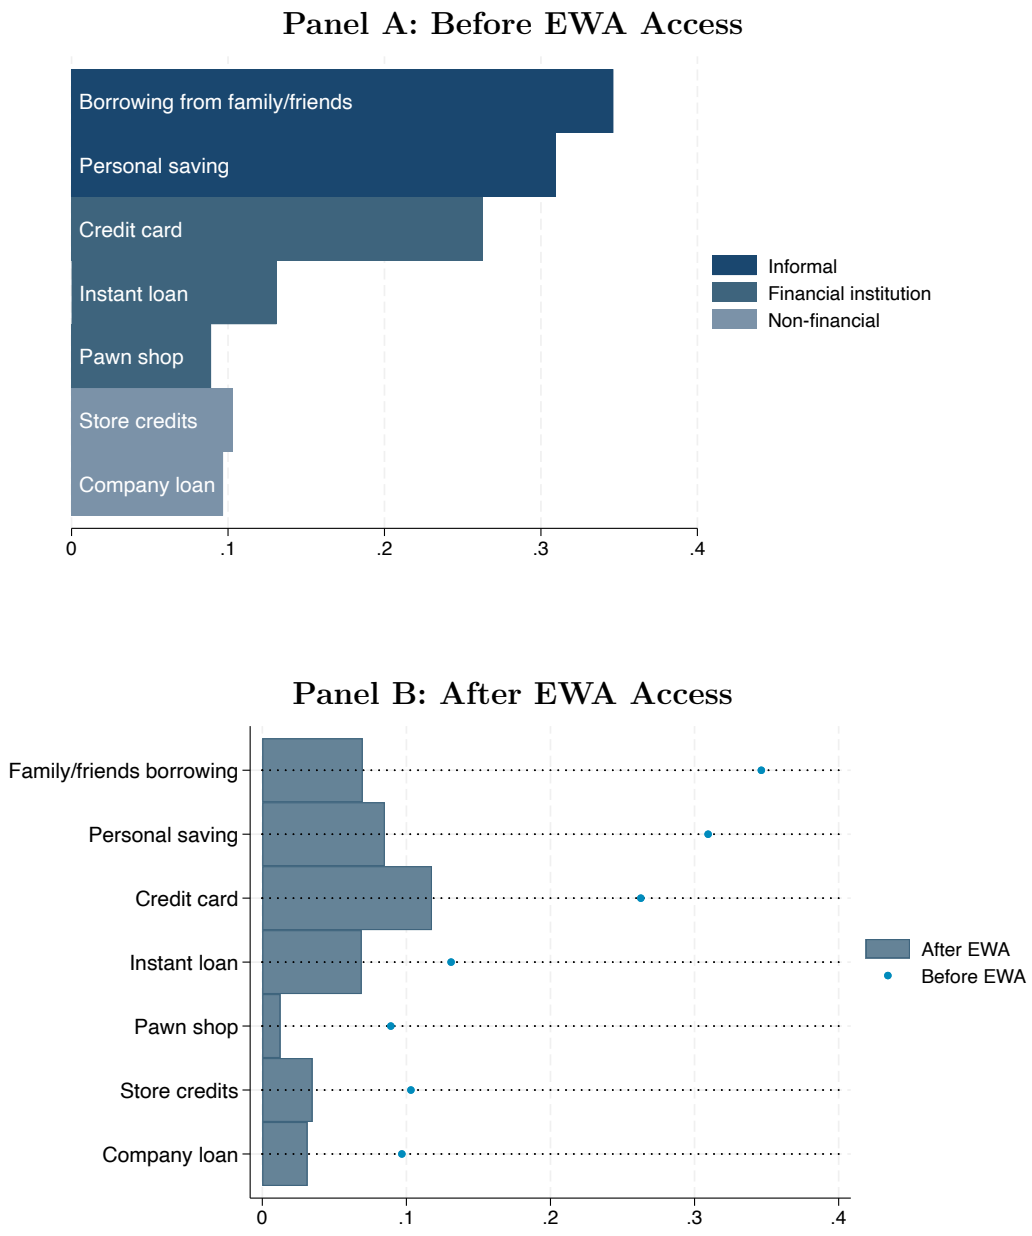


Figure 4: **Financial Products Used Before and After EWA Access**

Notes: These figures display the proportions of financial products that respondents report using for their main purposes of use. Panel A plots answers to the survey question "What are the financial products that you have used before Minu EWA was available to you for such purpose(s)?". Panel B plots the proportion of uses of financial products other than Minu EWA after it became available, among respondents who still use multiple financial products.

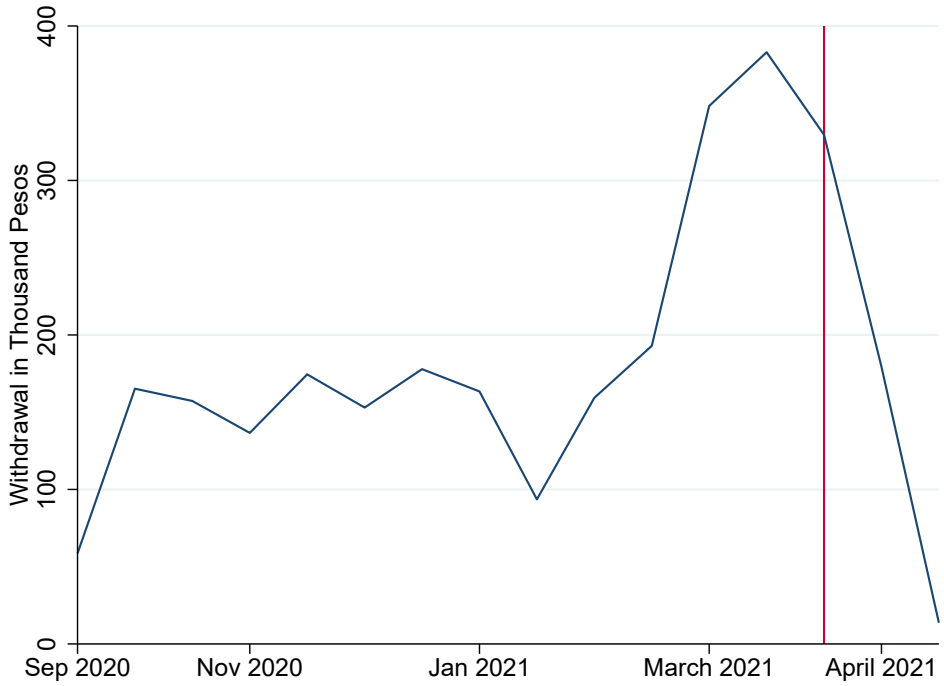


Figure 5: **EWA Withdrawals Around Semana Santa**

Notes: This figure plots the aggregate withdrawal amount on Minu in thousands of pesos from September 2020 to April 2021 in our sample. The red line indicates the week of "Semana Santa", a major holiday in Mexico.

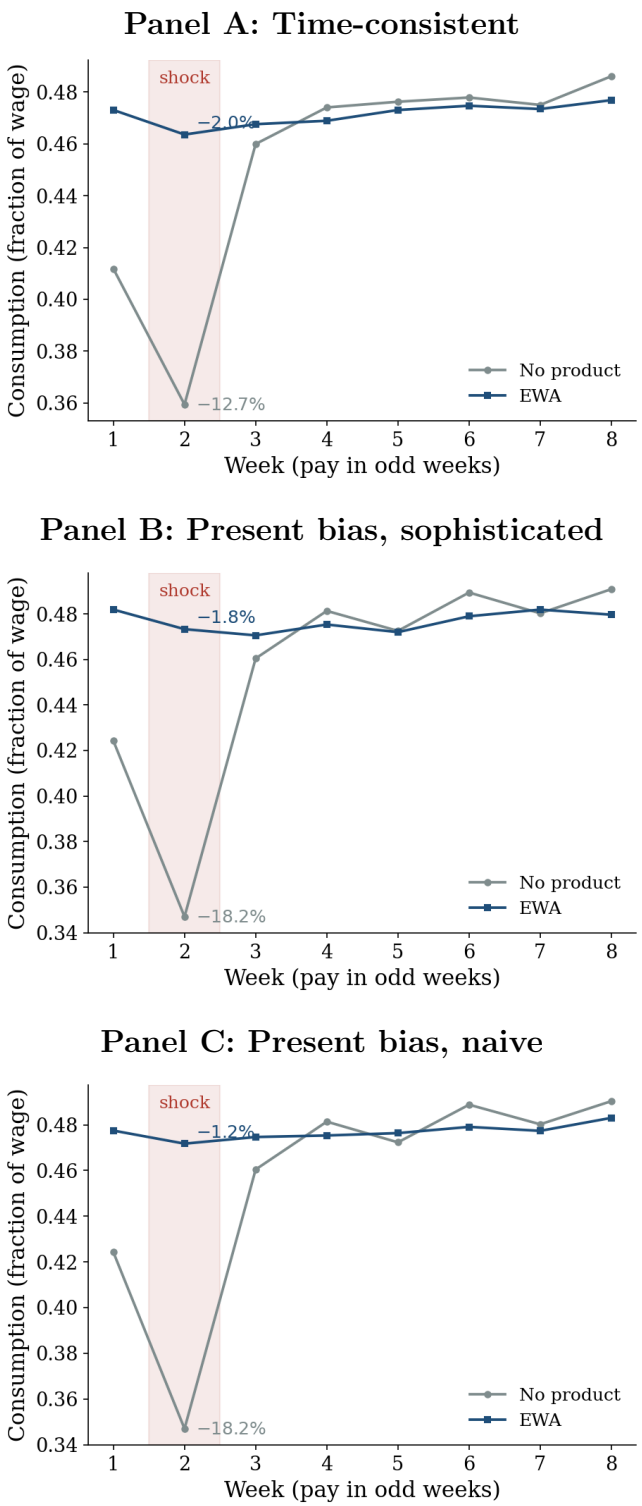


Figure 6: **Consumption Paths Around a Liquidity Shock**

Notes: This figure plots equilibrium weekly consumption over four pay cycles for a worker entering the first payday with cash-on-hand of 0.9 wages and facing an expenditure shock $\xi = 0.13$ in the first mid-cycle week (shaded), under the no-product and EWA regimes. Panel A displays the time-consistent type ($\beta = 1$), Panel B the sophisticated present-biased type ($\beta = 0.95$), and Panel C the naive present-biased type ($\beta = 0.95$), who believes each period that her future selves will behave in a time-consistent manner. The annotations report the drop in consumption in the shock week relative to the preceding payday week under each regime. Wages are paid in odd weeks. Parameters are as in Table 3.

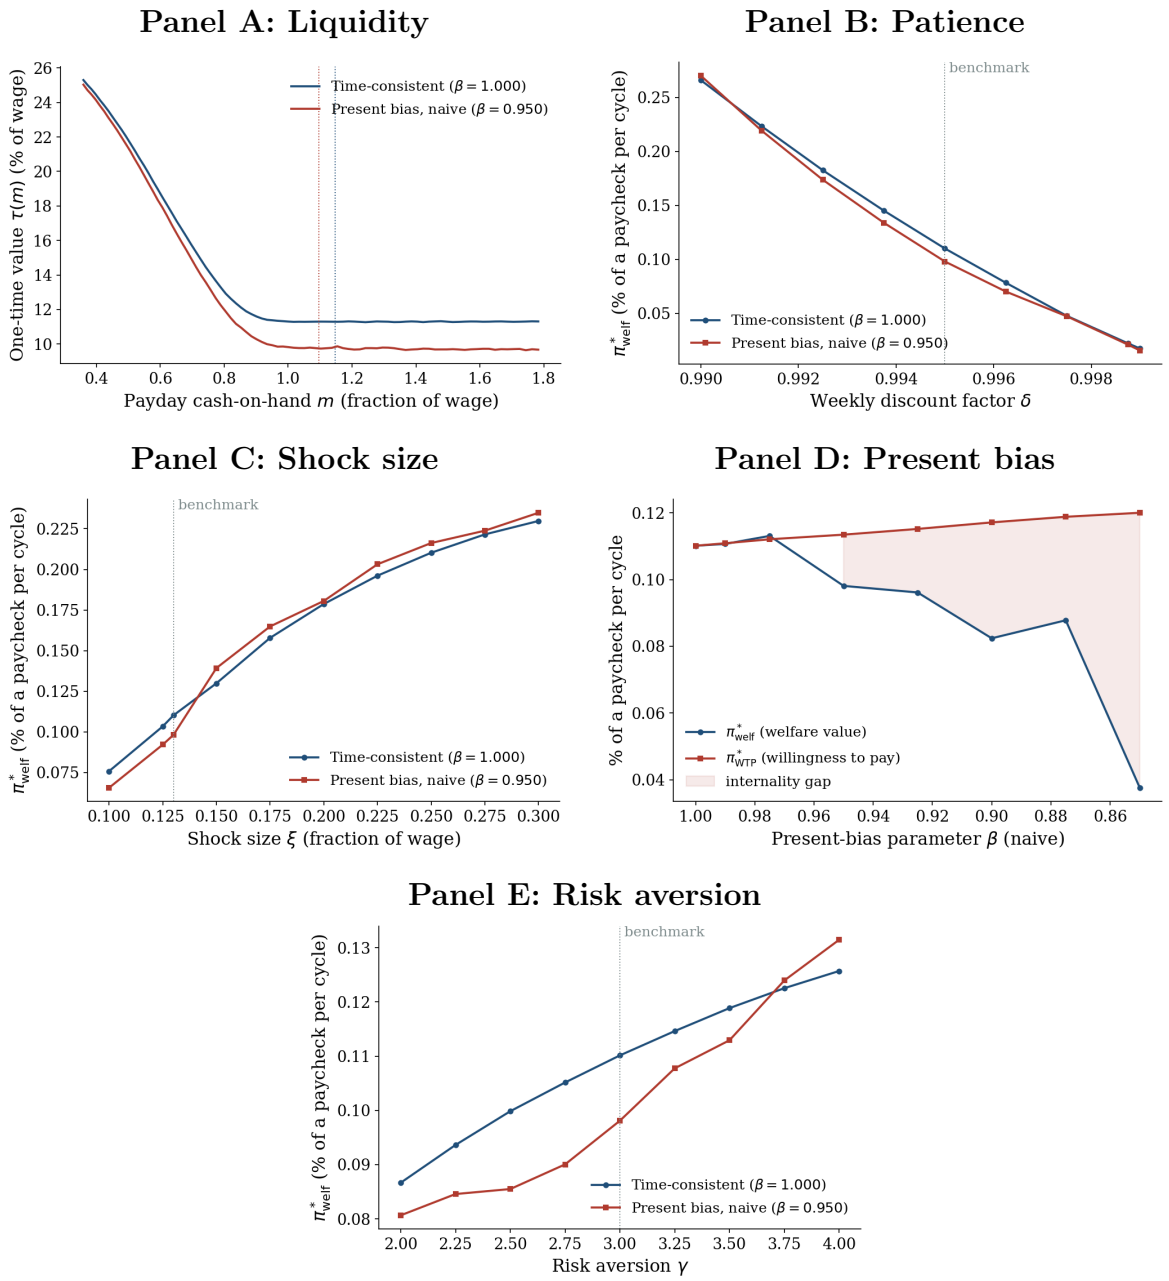


Figure 7: **For Whom is EWA Valuable?**

Notes: This figure reports comparative statics of the value of EWA access for the time-consistent type ($\beta = 1$) and the *naive* present-biased type ($\beta = 0.95$). The sophisticated type is not traced here, as her equilibrium is an intra-personal fixed point whose numerical stability restricts the range of bias that can be covered (see Appendix A). Each panel varies one dimension around the benchmark calibration of Table 3, marked by a dotted vertical line. Panel A plots the one-time equivalent value of permanent free access, $\tau(m)$, against payday cash-on-hand m, with the dotted lines indicating the ergodic mean cash-on-hand of each type in the no-EWA economy. Panels B, C and E plot the welfare-consistent access price π^*_{welf} against the weekly discount factor, the shock size and risk aversion. Panel D plots π^*_{welf} and the willingness to pay π^*_{WTP} of the naive type against β, the shaded area between them being the internality gap; its non-monotonicity near $\beta = 0.875$ reflects a discrete switch in equilibrium usage. All values are in percent of the semimonthly wage, per pay cycle in Panels B–E. Panels B–E evaluate indifference in expected value over the ergodic distribution of the no-product economy, which reproduces the access prices of Table 3 to within a few percent.

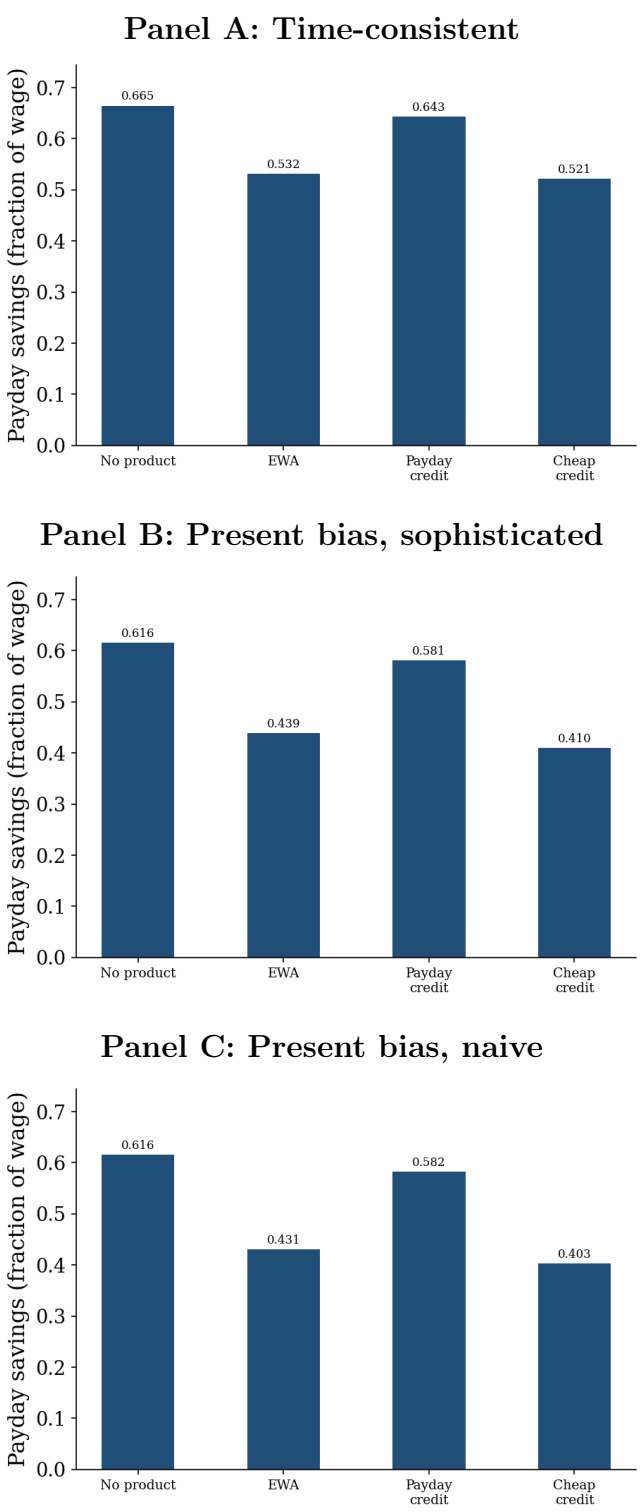


Figure 8: **Equilibrium Payday Savings: EWA versus Credit**

Notes: This figure plots ergodic average payday savings (in fractions of the semimonthly wage) under four regimes: no product, EWA, payday-style credit ($R^{bor} = 1.05$ per week, with a minimum loan size), and frictionless cheap credit ($R^{bor} = 1.02$ per week, no minimum size). Panel A displays the time-consistent type ($\beta = 1$), Panel B the sophisticated present-biased type ($\beta = 0.95$), and Panel C the naive present-biased type ($\beta = 0.95$). All other parameters are as in Table 3.

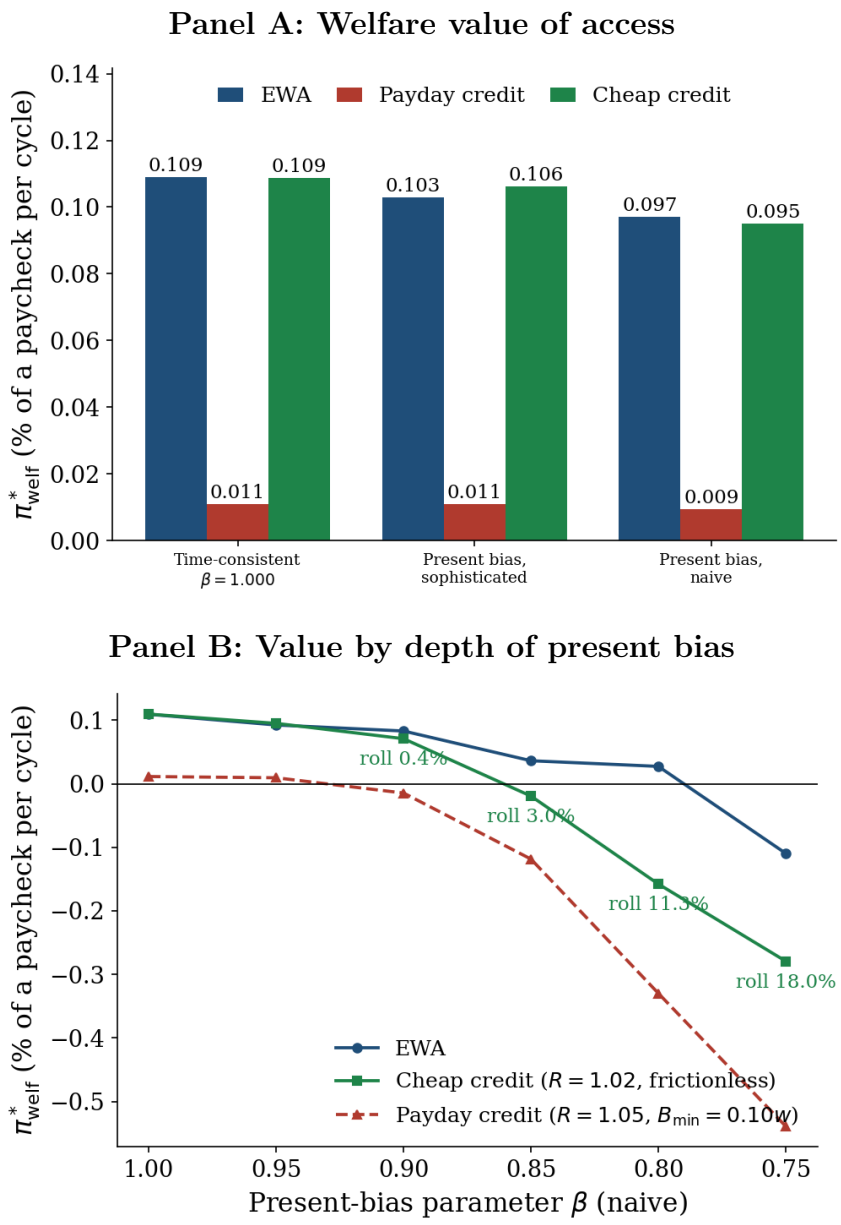


Figure 9: **Welfare Value of EWA versus Credit Access**

Notes: Panel A plots the welfare-consistent access price π^*_{welf} (in percent of the semimonthly wage, per pay cycle) for EWA, payday-style credit ($R^{bor} = 1.05$ per week, with a minimum loan size $B_{\min} = 0.10\,w$), and frictionless cheap credit ($R^{bor} = 1.02$ per week, no minimum size), for the time-consistent ($\beta = 1$), sophisticated present-biased ($\beta = 0.95$) and naive present-biased ($\beta = 0.95$) types. Panel B plots π^*_{welf} for the same three products for *naive* present-biased workers, who believe their future selves to be time-consistent, as a function of β down to 0.75; the annotations report the ergodic debt carried past payday under frictionless cheap credit, which is zero for EWA at every β by construction, since the advance is settled at the next payday. π^*_{welf} is the recurring access fee that leaves the long-run criterion indifferent between having and not having the product, evaluated at the ergodic mean cash-on-hand of the no-product economy. All other parameters are as in Table 3.

Tables

Table 1: **Summary Statistics**

Panel A. Employee Characteristics						
	EWA users (N = 3,779)			Non-EWA users (N = 40,262)		
	Mean	**p25**	**p75**	**Mean**	**p25**	**p75**
Age	32.3	26	37	33.1	25	39
Male	0.70	0	1	0.60	0	1
Start year	2018.2	2017	2019	2018.7	2018	2020
% Low-rank employees	36.5	-	-	59.3	-	-
% Mid-rank employees	34.4	-	-	27.5	-	-
% High-rank employees	29.1	-	-	13.2	-	-
% Paid semimonthly	70.4	-	-	71.5	-	-
Leaves Firm (1/0)	0.13	0	0	0.20	0	0

Panel B. EWA User Activity						
	Semimonthly (N = 12,889 pay cycles)			Weekly (N = 5,488 pay cycles)		
	Mean	**p25**	**p75**	**Mean**	**p25**	**p75**
Withdrawals per active cycle	1.35	1.00	1.00	1.29	1.00	1.00
Amount per active cycle (MXN)	953.2	430.0	1,000.0	385.1	150.0	440.0
Amount / own paycheck (%)	12.1	7.6	15.9	29.9	11.3	32.1
% of pay cycles active	57.8	30.0	88.9	31.8	8.3	50.0

Notes: This table reports summary statistics for the main variables used in the analysis. Panel A is at the employee level and covers the estimation sample of the retention analysis of Section 6, that is, every worker employed at some point between the introduction of EWA at her employer and the end of the observation window. It reports worker age, gender, start year, job rank, pay frequency, and the share of workers leaving the firm, separately for workers who ever used EWA and workers who never did. The three job-rank shares are computed among workers whose rank is classified; rank is unclassified for 29.5% of non-users and 0.0% of users, almost all of them company-1 workers for whom neither income nor a job title is recorded. Panel B is at the worker–pay cycle level and covers EWA users, split by pay frequency. Its first three rows condition on pay cycles in which the worker made at least one withdrawal; the last row is computed at the worker level, over 2,672 semimonthly and 1,121 weekly users. A semimonthly pay cycle spans about fifteen days and a weekly cycle seven, and the two groups draw very different paychecks, so the peso row is not comparable across the two columns whereas the paycheck-scaled row is. The shares reported here are averages across workers; the moments to which the framework of Section 5 is calibrated are computed at the worker–pay cycle level and are therefore slightly lower. Amounts are in Mexican pesos.

Table 2: **Earned Wage Access Adoption and Worker Characteristics**

	EWA user=1 (1)
Male=1	0.423***
	(0.004)
Young=1	0.325***
	(0.004)
Low rank=1	-0.136**
	(0.004)
Start two years or less before EWA introduction	0.246***
	(0.010)
Start after EWA introduction	-2.33***
	(0.010)
Observations	23,472
R^2	0.082

Notes: This table presents coefficients from a logistic regression of an indicator for using EWA on the worker's gender, age group, job rank, start-time group, and employer. Standard errors are clustered at the start-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: **Calibration, Model Fit, and the Value of EWA**

Panel A: Parameters

Parameter	Value	Source / target
Semimonthly wage w	1.0	Normalization
Withdrawal fee $1-\eta$	2%	Provider fee schedule
Withdrawal cap $\bar{x}$	$0.30\,w$	50% of earned wages, at the observed withdrawal timing
Minimum withdrawal $x_{\min}$	$0.015\,w$	Smallest withdrawal observed (≈100 pesos)
Risk aversion γ	3.0	Standard range
Weekly discount factor δ	0.995	Usage frequency; thin liquid buffers
Shock probability p	0.30	Usage frequency
Shock size ξ	$0.13\,w$	Share of withdrawals at the cap
Present bias β	$\{1.0,\ 0.95\}$	Frequent-user usage patterns
Taste-shock scale σ	3×10^{-4}	Numerical smoothing; see Appendix A

Panel B: Usage moments (per user–pay cycle, conditional statistics in %)

	Data	$\beta=1.0$	$\beta=0.95$ (naive)	$\beta=0.95$ (soph.)
Share of cycles with a withdrawal	48.4 (0.70)	40.0	79.7	72.8
Mean withdrawal | use (% of wage)	12.4 (0.10)	9.7	12.3	12.8
Utilization of the cap | use	44.1 (0.35)	32.5	40.9	42.6
Share of withdrawals at cap	0.7 (0.12)	0.7	2.9	2.6
Shock-driven share of usage (untargeted)	≈ 50 (survey)	56.3	37.4	40.6

Panel C: The value of EWA (% of the semimonthly wage)

	$\beta=1.0$	$\beta=0.95$ (naive)	$\beta=0.95$ (soph.)
Welfare value π^*_{welf} (per cycle)	0.109	0.097	0.103
Willingness to pay π^*_{WTP} (per cycle)	0.109	0.112	0.104
Internality gap $\pi^*_{\text{WTP}}-\pi^*_{\text{welf}}$	0.000	0.015	0.001
One-time equivalent $\tau(\bar{m})$	11.3	9.7	10.3
Expected fees paid (per cycle)	0.078	0.196	0.186
Payday buffer: no EWA → EWA	0.67 → 0.53	0.62 → 0.43	0.61 → 0.44

Notes: Panel A reports the calibrated parameters of the stationary model. Panel B compares equilibrium usage in the model with the usage moments documented in Section 3: the share of user–pay cycles with a withdrawal and the mean withdrawal conditional on use are from Figures 1 and 2; the time-consistent type is calibrated to the average user and the present-biased type is compared with the frequent-user group of Section 3; the shock-driven share of usage is untargeted and compared with the survey evidence on purposes of use (Figure 3). All moments are computed on workers paid semimonthly. Bootstrap standard errors, clustered on the worker, are in parentheses. The naive type is the same present-biased type under naivete about her own bias: her willingness to pay then reflects usage plans she does not follow, opening a positive internality gap. The sophisticated type correctly anticipates her own future behavior. Panel C reports the per-cycle access prices defined in Section 5, the one-time equivalent value of permanent free access evaluated at the ergodic mean cash-on-hand $\bar{m}$ of the no-EWA economy, expected equilibrium withdrawal fees, and ergodic mean payday savings without and with EWA. All monetary quantities are expressed in percent of the semimonthly wage. Every column is solved on the same grids and at the same taste-shock scale: $\sigma=3\times10^{-4}$, 181 points for cash-on-hand and for each carryover choice, and 51 points for the withdrawal, one at zero and fifty spanning $[x_{\min},\bar{x}]$. The internality gap of the naive type is a lower bound: it widens to about a fifth of π^*_{welf} as σ is taken to zero, while the time-consistent values and all usage moments are unaffected (Appendix A).

Table 4: **EWA versus Credit Access**

	Time-consistent ($\beta = 1.0$)			Present-biased, sophisticated ($\beta = 0.95$)			Present-biased, naive ($\beta = 0.95$)		
	EWA	Payday credit	Cheap credit	EWA	Payday credit	Cheap credit	EWA	Payday credit	Cheap credit
Welfare value π^*_{welf} (% wage/cycle)	0.109	0.011	0.109	0.103	0.011	0.106	0.097	0.009	0.095
Share of cycles with use (%)	40.0	4.7	46.9	72.8	10.5	80.5	79.7	10.3	81.9
Mid-cycle position (% of wage)	−3.9	−0.5	−4.4	−9.3	−1.3	−11.4	−9.8	−1.3	−12.4
Debt carried past payday (% of wage)	0	0.000	0.000	0	0.000	0.009	0	0.000	0.009
Fees or interest paid (% wage/cycle)	0.078	0.027	0.087	0.186	0.065	0.229	0.196	0.064	0.248

Notes: This table compares EWA with two credit regimes that replace it in an otherwise identical environment. Payday-style credit carries a borrowing rate of $R^{bor} = 1.05$ per week, in line with observed payday and informal lending rates, a minimum loan size $B_{\min} = 0.10w$ reflecting fixed origination costs. Frictionless cheap credit carries $R^{bor} = 1.02$ per week, matching EWA's effective cost, with no minimum size, isolating the role of the automatic repayment structure. π^*_{welf} is the recurring access fee that leaves the long-run criterion indifferent to the no-product baseline, evaluated at the ergodic mean cash-on-hand of the no-product economy. Mid-cycle positions are ergodic averages; for EWA they equal the outstanding advance (repaid automatically at the next payday), and for credit regimes net debt. Debt carried past payday is the ergodic average debt outstanding after wage receipt. The EWA columns are the solution reported in Table 3, computed with the minimum withdrawal $x_{\min} = 0.015w$ and the taste-shock scale $\sigma = 3 \times 10^{-4}$; the credit columns are computed with neither, since a minimum withdrawal is not a feature of a credit contract and the smoothing is not required where the borrowing problem has a unique solution. All other parameters are as in Table 3.

Table 5: **EWA Use and Turnover: Hazard Model**

	Leaves firm=1			
	(1)	(2)	(3)	(4)
Used Minu	-0.103**	-0.115**	-0.142***	-0.128***
	(0.046)	(0.046)	(0.048)	(0.047)
Male			0.083***	0.070***
			(0.026)	(0.026)
Age				-0.002
				(0.002)
Company x Frequency Strata	Y	Y	Y	Y
Start Quarter Strata	N	Y	Y	Y
Job Rank	N	N	strata	strata
Number of Employees	43,445	43,445	31,712	26,966
Observations	411,859	411,859	298,381	251,690

Notes: This table presents coefficients (log hazard ratios) from Cox proportional hazard regressions relating the hazard of a worker leaving their firm to an indicator for EWA usage; the implied proportional change in the hazard is $e^{\hat{\beta}} - 1$. Baseline hazards are stratified by company × pay frequency and, from column 2 on, by starting quarter, with job rank added as strata in columns 3 and 4. The remaining controls are the worker's gender (columns 3 and 4) and age (column 4). Standard errors are clustered by employee. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: **Hazard Model: Heterogeneity**

	Leaves firm=1					
	Gender		Job Rank		Age	
	Female	Male	Low	Mid-to-High	Young	Older
	(1)	(2)	(3)	(4)	(5)	(6)
Used Minu	-0.194**	-0.083	-0.140*	-0.134**	-0.215***	0.000
	(0.081)	(0.057)	(0.073)	(0.061)	(0.064)	(0.068)
Company x Frequency Strata	Y	Y	Y	Y	Y	Y
Start Quarter Strata	Y	Y	Y	Y	Y	Y
Age	Y	Y	Y	Y	N	N
Gender	N	N	Y	Y	Y	Y
Rank	Y	Y	N	N	Y	Y
Number of Employees	9,951	17,015	14,839	12,127	14,747	12,219
Observations	92,024	159,666	133,306	118,384	134,014	117,676

Notes: This table presents coefficients (log hazard ratios) from Cox proportional hazard regressions relating the hazard of a worker leaving their firm to an indicator for EWA usage. The sample is split on worker gender (columns 1 and 2), worker job rank (columns 3 and 4), and worker age (columns 5 and 6), where young workers are those born after 1989. Baseline hazards are stratified by company $\times$ pay frequency and by starting quarter throughout. The set of controls includes the worker's gender, age, and job rank, each where it is not the splitting variable. All columns condition on workers whose rank is classified. Standard errors are clustered by employee. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: **EWA Usage Propensity and Turnover**

	Exit firm within:			
	2 months	4 months	6 months	8 months
Post × High Propensity To Adopt	-0.043**	-0.051*	-0.051	-0.058*
	(0.019)	(0.027)	(0.031)	(0.033)
Post	-0.044**	-0.084**	-0.131***	-0.185***
	(0.021)	(0.035)	(0.042)	(0.044)
High Propensity To Adopt	0.026	0.033	0.040	0.039
	(0.019)	(0.026)	(0.028)	(0.031)
Company FE	Y	Y	Y	Y
Company × Post	Y	Y	Y	Y
Cluster level	Start quarter	Start quarter	Start quarter	Start quarter
Observations	20,968	20,968	20,968	20,968

Notes: This table presents estimates from the propensity-based difference-in-differences design. The DiD specification regresses an indicator for leaving the firm within a given horizon (2, 4, 6, and 8 months) on the interaction of a post-EWA-introduction indicator with the High Propensity group indicator, along with company fixed effects and company-by-post fixed effects to absorb time-varying firm-level shocks. Standard errors are clustered at the start-quarter level to account for within-cohort correlation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A Appendix

A.1 Proof of Proposition 1

Consider the mid-cycle stage with no shock, liquid resources s, and EWA access. The mid-cycle self chooses consumption c_B, carryover $k \geq 0$, and withdrawal $x \in [0, \bar{x}]$ subject to $c_B = s + \eta x - k$, to maximize

$$u(c_B) + \beta\delta\, \mathbb{E}\left[V_A(k + w - x - \pi)\right],$$

where $V_A(\cdot)$ is the (exponentially discounted) continuation value of equilibrium behavior from the next payday onward, which is strictly increasing in cash-on-hand since marginal utility is strictly positive.

Suppose, contrary to the claim, that the optimal policy has $k^* > 0$ and $x^* > 0$. Consider the perturbation that lowers the withdrawal to $x^* - \varepsilon$ and the carryover to $k^* - \eta\varepsilon$, for $\varepsilon \in (0, \min\{x^*, k^*/\eta\})$. Consumption is unchanged: $c_B = s + \eta(x^* - \varepsilon) - (k^* - \eta\varepsilon) = c_B^*$. Cash-on-hand at the next payday, $m' = k + w - x - \pi$, rises by $(1 - \eta)\varepsilon > 0$. Since V_A is strictly increasing, the perturbation strictly raises the mid-cycle self's objective, a contradiction. Hence $x^* = 0$. ■

Four remarks are in order. First, the argument uses only the monotonicity of the continuation value: it requires no differentiability or concavity of V_A, properties that are not guaranteed in Markov-perfect equilibria of quasi-hyperbolic consumers. When V_A is differentiable, the same conclusion follows from first-order conditions: interior savings imply $u'(c_B^*) = \beta\delta\, \mathbb{E}[V_A'(k^* + w - x^* - \pi)]$, so the marginal value of a withdrawal is $\eta\, u'(c_B^*) - \beta\delta\, \mathbb{E}[V_A'] = (\eta - 1)\, u'(c_B^*) < 0$. Second, the present-bias parameter β plays no role in the argument, so the result holds for any $\beta \in (0, 1]$; likewise, nothing depends on the shock realization, which only shifts the mid-cycle resources s. In any mid-cycle state with interior carryover the optimal withdrawal is therefore zero, so all equilibrium usage occurs at hand-to-mouth states, which shocks make more likely to be reached. Third, the argument fails exactly when $k^* = 0$ binds: then $u'(c_B^*) > \beta\delta\, \mathbb{E}[V_A']$ is possible, and a strictly positive withdrawal can be optimal, characterized by $\eta\, u'(s + \eta x^*) = \beta\delta\, \mathbb{E}[V_A'(w - x^* - \pi)]$ when interior. EWA usage without a shock therefore reveals a binding mid-cycle liquidity constraint. Fourth, because the environment is stationary, the proposition applies

to every pay cycle identically; no restriction on usage needs to be imposed anywhere in the analysis, in contrast with finite-horizon settings in which terminal-date borrowing motives contaminate non-shock usage.

A.2 Numerical Solution

We solve for a stationary Markov-perfect equilibrium of the intra-personal game on a discretized state space. The state at the payday stage is cash-on-hand m, on a grid of 181 points concentrated at low liquidity; carryover choices lie on grids of equal size, and EWA withdrawals on a grid of 51 points, one at zero and fifty spanning $[x_{\min}, \bar{x}]$. Given a candidate continuation value V_A on the payday grid, we solve the mid-cycle problem over withdrawals and carryovers for each savings level and shock realization, with current choices evaluated under $\beta\delta$ weighting and continuation values interpolated linearly; we then solve the payday problem analogously. Rather than taking the maximum over the discretized choice set, we let each alternative carry an independent extreme-value taste shock of scale σ, so that stage values are smooth in V_A and behavior at each state is a distribution over choices rather than a point (Rust, 1987, Iskhakov et al., 2017). We solve at a decreasing sequence of σ, from 3×10^{-2} down to 3×10^{-4}, starting each solution from the preceding one; without this the intra-personal fixed point of the sophisticated type does not converge at the degrees of bias we report. The minimum withdrawal size matters here beyond its institutional role: by bounding the smallest available advance away from zero it keeps the taste shock from manufacturing withdrawals that differ arbitrarily little from no withdrawal at all, which is what makes the reported moments invariant to the fineness of the withdrawal grid. The exponentially discounted value of the implied equilibrium behavior is computed exactly by solving the linear system associated with the policy-induced transition operator over the pay cycle (a form of policy iteration with exact evaluation), and the procedure iterates until the value function converges, with mild damping to ensure convergence of the intra-personal fixed point. The ergodic distribution of cash-on-hand is obtained from the policy-induced transition matrix, and all reported moments are computed under this distribution. In the credit regimes, the state space is extended to negative positions, with interest applied weekly to negative balances, the flow disutility of debt subtracted from period payoffs, and the minimum loan size and borrowing limit imposed as restrictions on the carryover grid. Access prices π^* are found by solving the model over a sequence of fees

and locating the indifference point via a secant iteration on the (nearly linear) map from fee to value; welfare comparisons across regimes use identical grids on the common (positive) region of the state space to avoid discretization bias. For naive workers, the perceived continuation value is the time-consistent ($\beta = 1$) solution of the same regime; actual policies at each stage maximize the current, β-weighted objective against this perceived continuation, and welfare is computed by exact policy evaluation of the resulting actual behavior under exponential discounting. Because actual behavior is a single best response to a fixed perceived value function, the naive solution involves no intra-personal fixed point and is numerically stable at any depth of bias. Finally, in the timing-only economy ($p = 0$) used in the decomposition of Section 5, the deterministic game admits a second, no-use stationary point for the present-biased type, supported by a kink in the discretized continuation value (Krusell and Smith, 2003); because the usage equilibrium is the unique outcome of the solution algorithm at every $p > 0$ we examine (down to $p = 0.001$), we impose it as the $p \to 0$ selection by warm-starting from the no-EWA value function. The comparative statics of Figure 7 are computed under the same selection convention, with indifference evaluated in expected value over the ergodic distribution of the no-product economy, which smooths the discrete regime switches that evaluation at a single state amplifies; at the benchmark, the resulting access prices coincide with those of Table 3 to within a few percent. Because those comparative statics use the naive type, whose problem involves no intra-personal fixed point, the range of bias they cover is not restricted by numerical stability. Finally, the access prices of a present-biased worker depend on σ in a way those of a time-consistent worker do not. Perturbing the choices of an agent who is optimizing is second order for her own criterion, so the willingness to pay is insensitive to σ and every time-consistent price moves by 0.2% across three orders of magnitude in σ; it is first order for a criterion the agent is not optimizing, which is what the long-run welfare value is for a biased worker. Randomizing her choices therefore undoes part of her bias, and the internality gap of the naive type, 16% of π^*_{welf} at the calibrated bias and at $\sigma = 3 \times 10^{-4}$, rises to 20% as σ is taken to zero. We report the former and record the latter as the bound.

A.3 Additional Figures

This appendix collects figures referenced in the main text but placed here for space. Figure A.1 plots the predicted probability of EWA usage across worker subgroups from the adoption logit; Figure A.2 reproduces the usage dynamics of Figure 2 for weekly-paid workers; Figures A.3 and A.4 document the heterogeneity of those dynamics across job ranks, for semimonthly- and for weekly-paid workers respectively; Figure A.6 reports how the reported purpose of use varies with the withdrawal amount; Figure A.7 reports workers' ranking of EWA among company benefits; and Figure A.8 reports the equilibrium EWA withdrawal policies of the model.

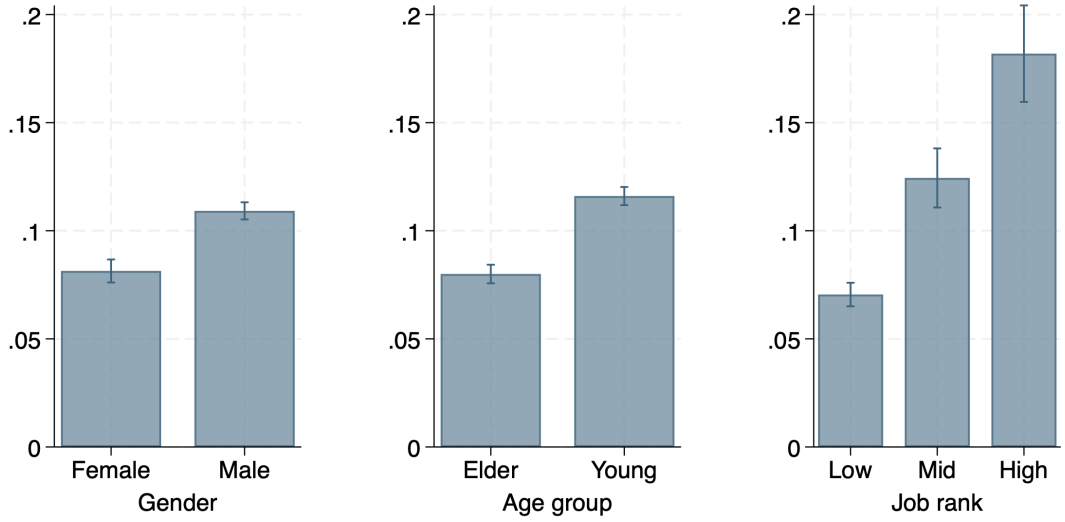


Figure A.1: **Predicted Probability of EWA Usage by Group**

Notes: This figure displays the predicted probability of EWA usage from the logit regressions for sub-groups of employees. For each subgroup, the predicted probability is computed using average covariates.

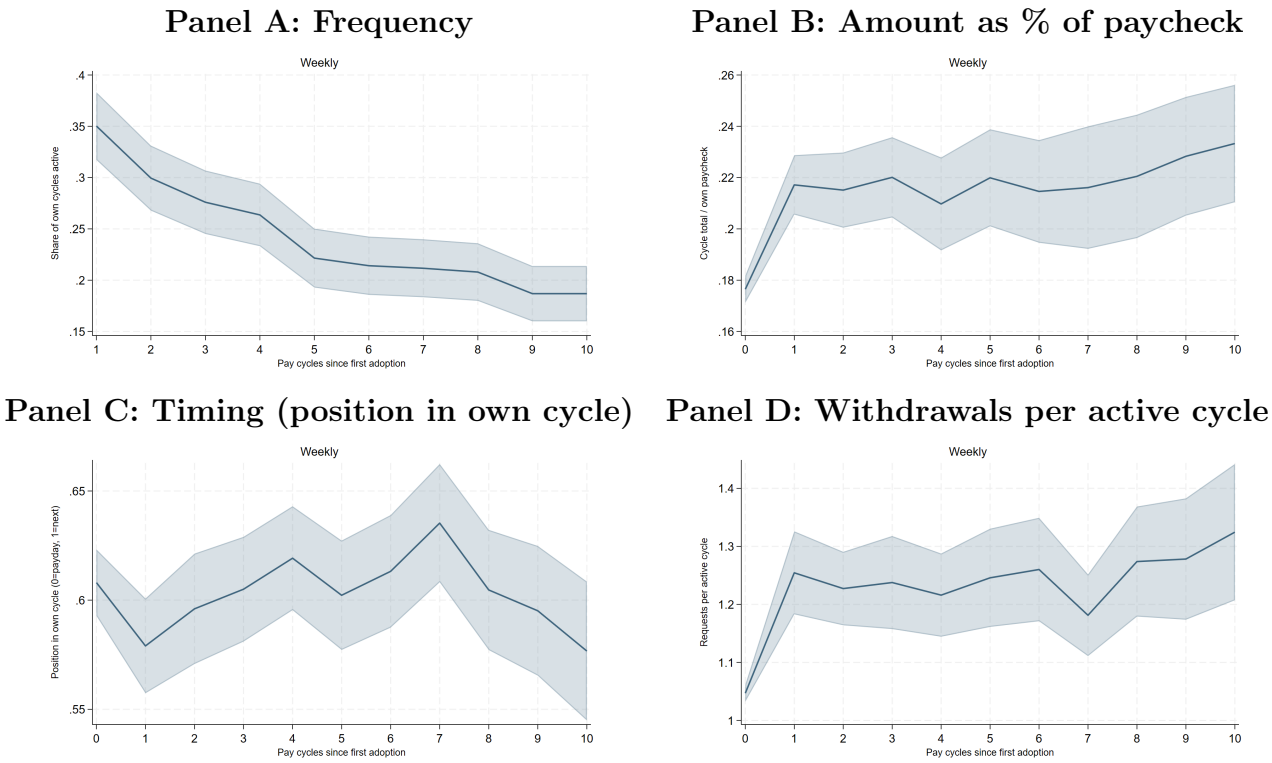


Figure A.2: **Average Withdrawal Behaviors over Time, Weekly-Paid Workers**

Notes: These figures reproduce Figure 2 for EWA users paid weekly. The horizontal axis counts a worker's own pay cycles since her first withdrawal, so it is a weekly rather than a semimonthly grid and the levels are not directly comparable with those of Figure 2. The panel is balanced: it retains the users who are observable for at least ten pay cycles after adoption. Frequency is the share of a worker's own cycles in which she makes at least one withdrawal; the three remaining panels condition on cycles with at least one withdrawal and report, respectively, the cycle total scaled by the worker's own paycheck, the position of the first request within the cycle (0 = payday, 1 = the next payday), and the number of requests per active cycle. Shaded bands are 95% confidence intervals from standard errors clustered on the worker.

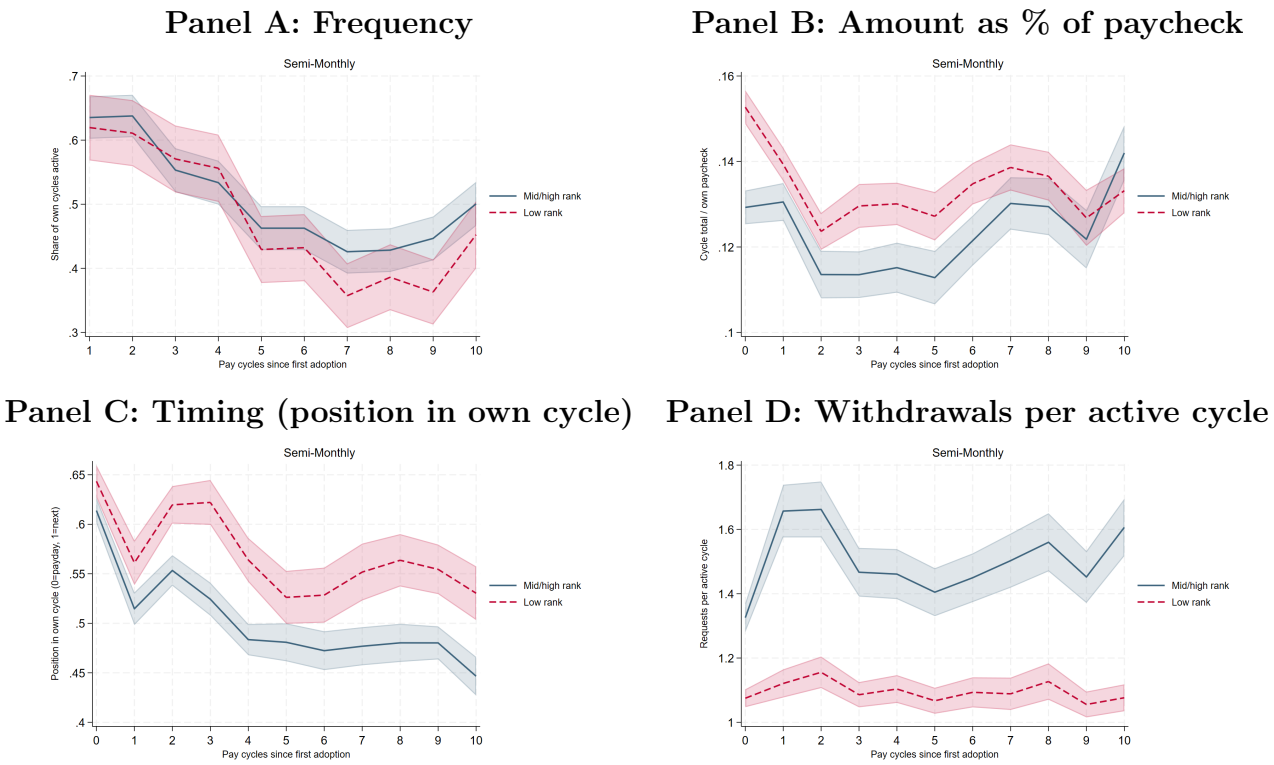


Figure A.3: **Withdrawal Behaviors over Time by Job Rank, Semimonthly-Paid Workers**

Notes: These figures present the evolution of the frequency, amount, timing and intensity of withdrawals over the pay cycles following adoption, for EWA users paid semimonthly, split by whether they hold a low-rank or a mid/high-rank position. The horizontal axis counts a worker's own pay cycles since her first withdrawal. The panel is balanced: it retains the users who are observable for at least ten pay cycles after adoption. Frequency is the share of a worker's own cycles in which she makes at least one withdrawal; the three remaining panels condition on cycles with at least one withdrawal and report, respectively, the cycle total scaled by the worker's own paycheck, the position of the first request within the cycle (0 = payday, 1 = the next payday), and the number of requests per active cycle. Shaded bands are 95% confidence intervals from standard errors clustered on the worker.

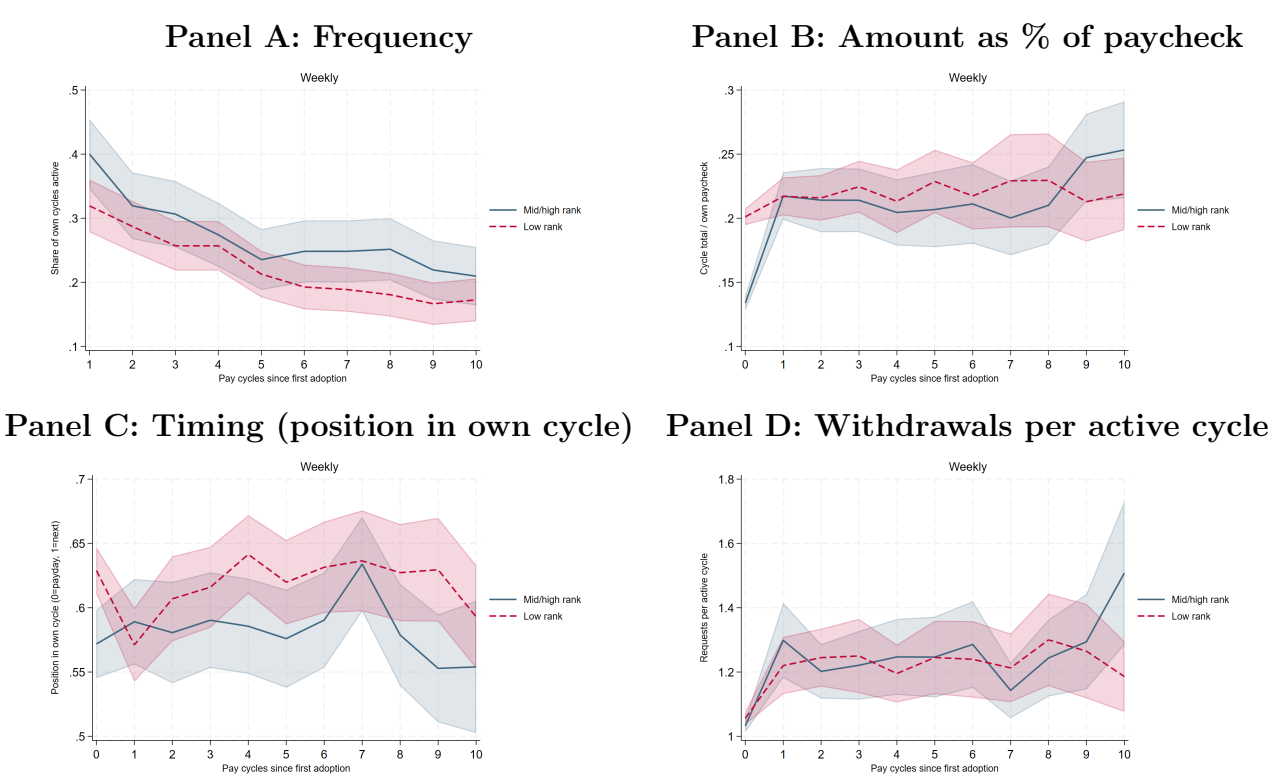


Figure A.4: **Withdrawal Behaviors over Time by Job Rank, Weekly-Paid Workers**

Notes: These figures present the evolution of the frequency, amount, timing and intensity of withdrawals over the pay cycles following adoption, for EWA users paid weekly, split by whether they hold a low-rank or a mid/high-rank position. The horizontal axis counts a worker's own pay cycles since her first withdrawal. The axis is a weekly rather than a semimonthly grid, so the levels are not directly comparable with those of Figure A.3. The panel is balanced: it retains the users who are observable for at least ten pay cycles after adoption. Frequency is the share of a worker's own cycles in which she makes at least one withdrawal; the three remaining panels condition on cycles with at least one withdrawal and report, respectively, the cycle total scaled by the worker's own paycheck, the position of the first request within the cycle (0 = payday, 1 = the next payday), and the number of requests per active cycle. Shaded bands are 95% confidence intervals from standard errors clustered on the worker.

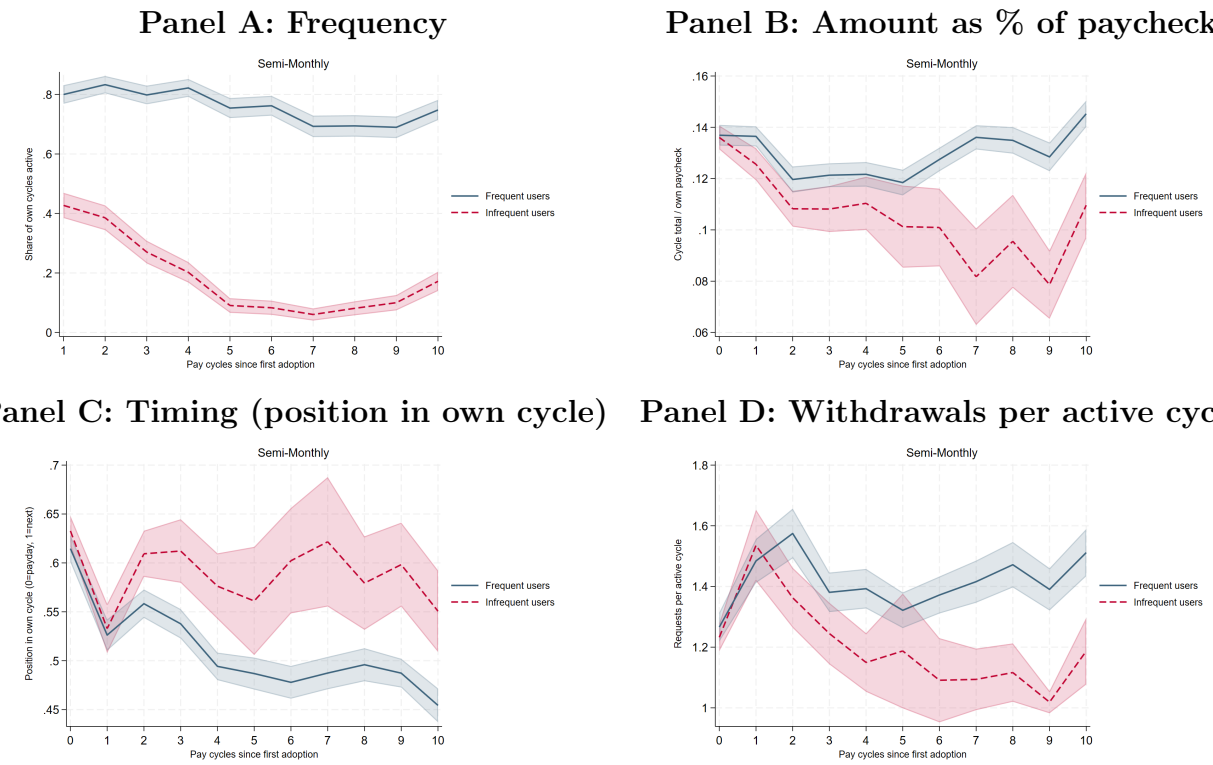
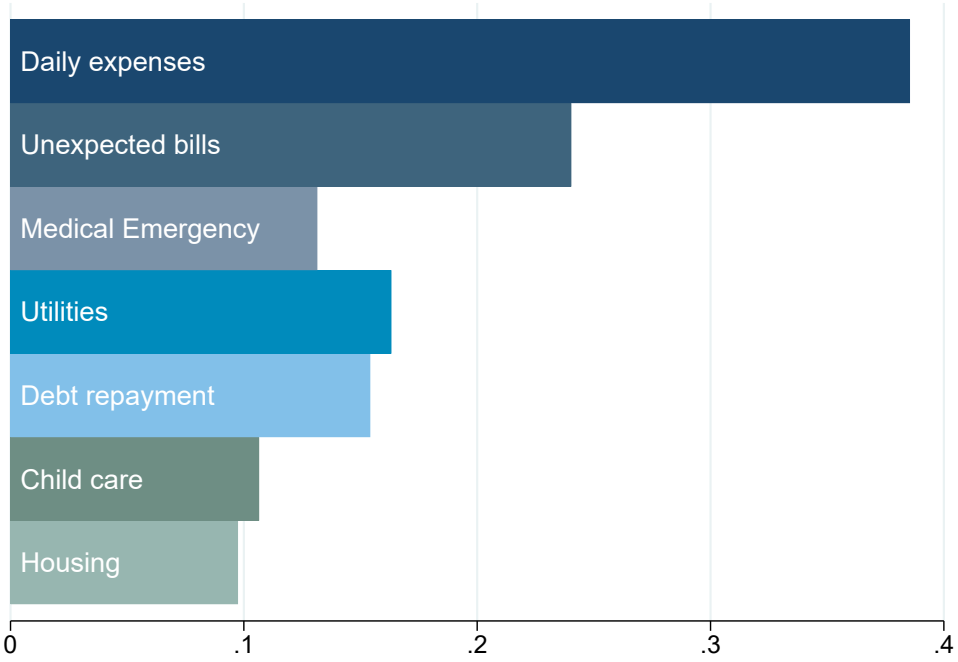


Figure A.5: **Withdrawal Behaviors over Time by Frequency, Semimonthly-Paid Workers**

Notes: These figures present the evolution of the frequency, amount, timing and intensity of withdrawals over the pay cycles following adoption, for EWA users paid semimonthly, split by whether they used more than half of the pay cycles since their first adoption. The horizontal axis counts a worker's own pay cycles since her first withdrawal. The panel is balanced: it retains the users who are observable for at least ten pay cycles after adoption. Frequency is the share of a worker's own cycles in which she makes at least one withdrawal; the three remaining panels condition on cycles with at least one withdrawal and report, respectively, the cycle total scaled by the worker's own paycheck, the position of the first request within the cycle (0 = payday, 1 = the next payday), and the number of requests per active cycle. Shaded bands are 95% confidence intervals from standard errors clustered on the worker.

Panel A: Less than or exactly as needed



Panel B: Maximum allowed

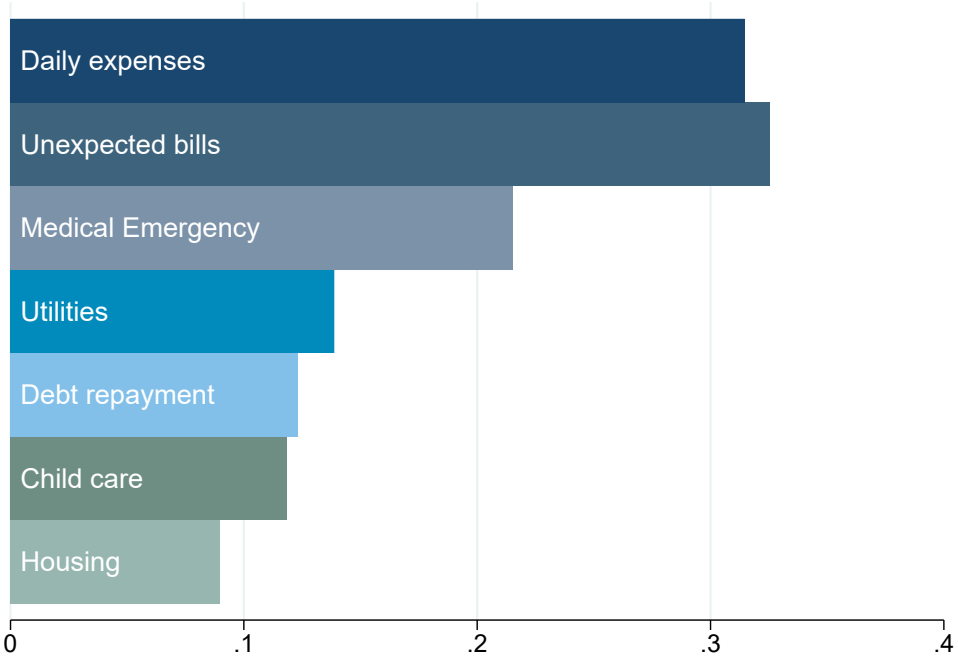


Figure A.6: **Proceeds Usage by Amount Withdrawn**

Notes: These figures display the reported uses of EWA proceeds by amount withdrawn, based on the survey question "How do you choose the amount you withdraw from Minu EWA when you use it?". Panel A plots the uses reported by respondents who withdraw less than or exactly as much as needed, and Panel B those reported by respondents who withdraw the maximum allowed.

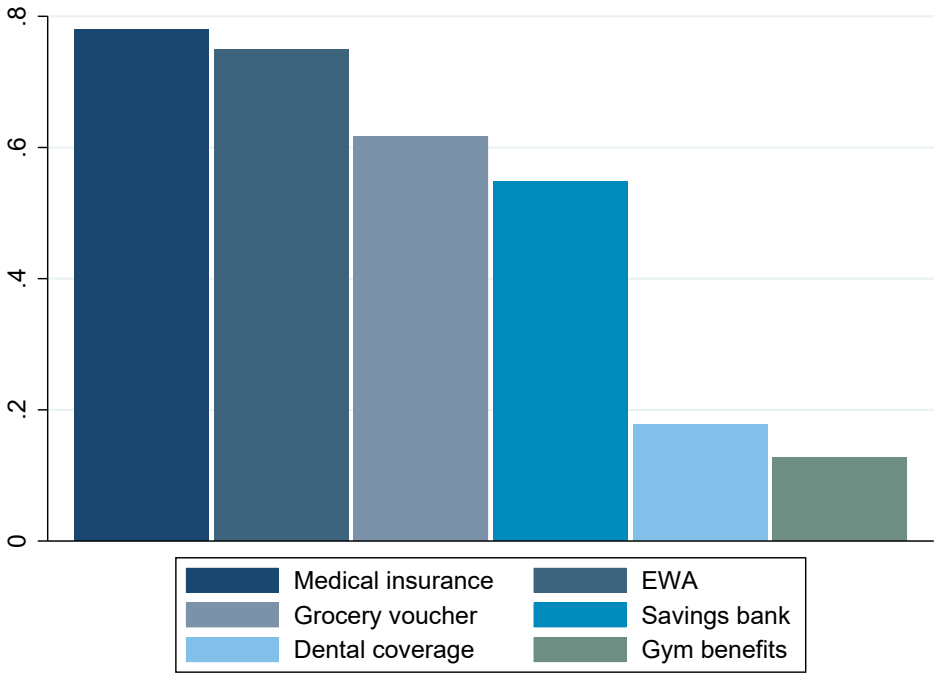


Figure A.7: **Preferences for Company Benefits**

Notes: This figure plots the proportion of respondents ranking each company benefit among their most valued benefits.

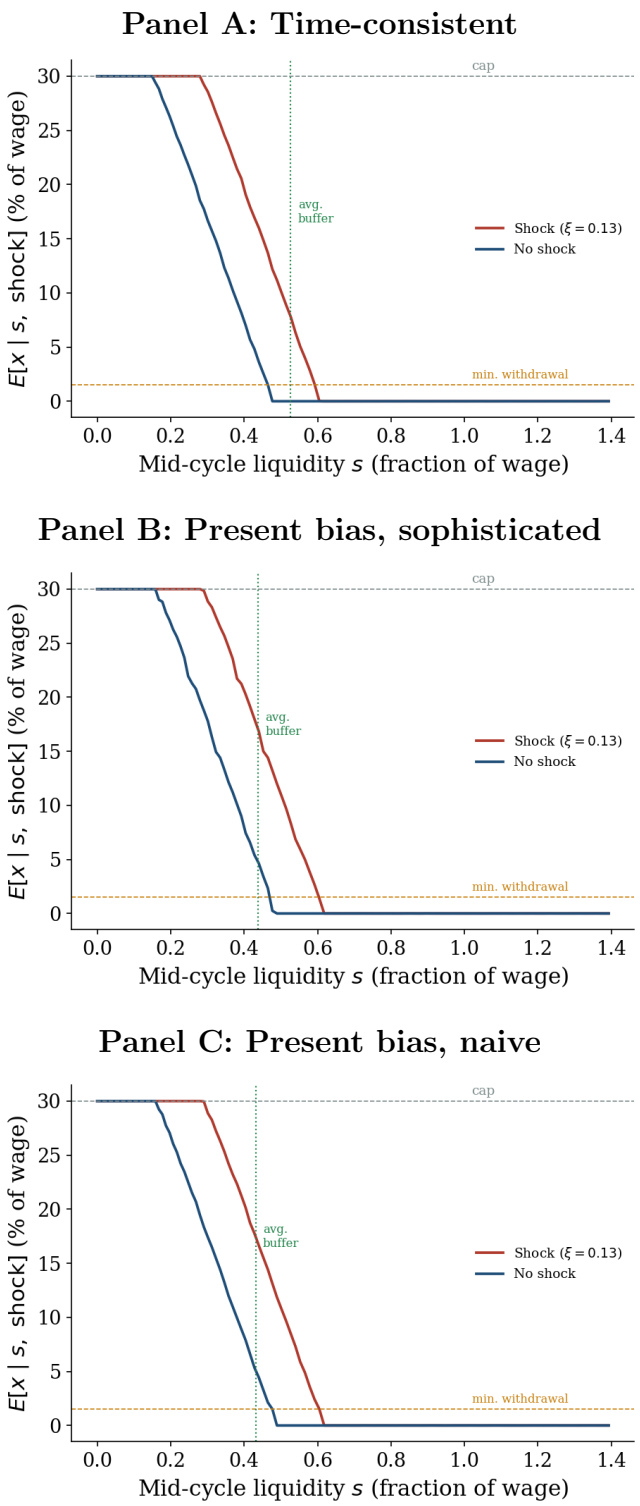


Figure A.8: **Equilibrium EWA Withdrawal Policies**

Notes: This figure plots the expected equilibrium EWA withdrawal $E[x \mid s]$ (in percent of the semimonthly wage) as a function of mid-cycle liquidity s, with and without a realized expenditure shock. Panel A displays the time-consistent type ($\beta = 1$), Panel B the sophisticated present-biased type ($\beta = 0.95$), and Panel C the naive present-biased type ($\beta = 0.95$). The upper dashed horizontal line indicates the contractual cap ($\bar{x} = 30\%$ of the semimonthly wage, i.e., 50% of earned wages at the observed withdrawal timing) and the lower one the minimum withdrawal ($x_{\min} = 1.5\%$ of the semimonthly wage); the dotted vertical line marks the ergodic mean payday savings of each type. Parameters are as in Table 3.

The remaining figures report the unconditional retention comparisons discussed in Section 6. As noted there, they are subject to a mechanical (immortal-time) bias and are shown for completeness only; the substantive results rely on the hazard and propensity-based designs.

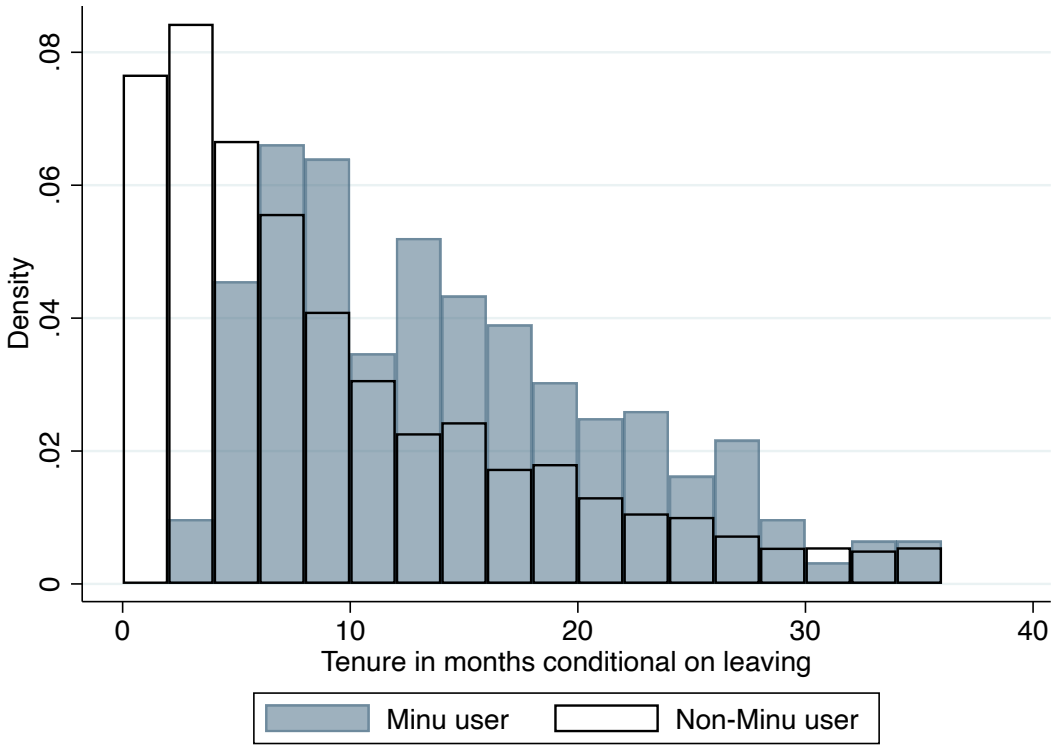
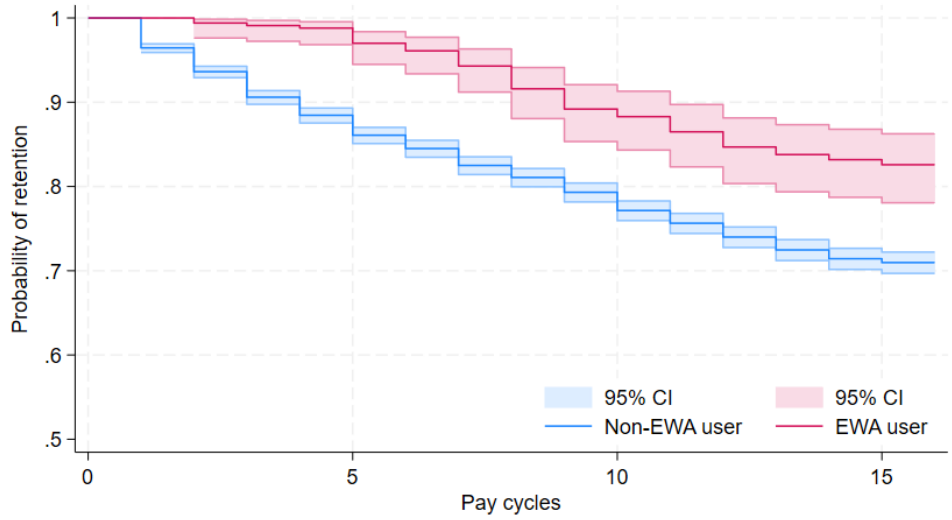


Figure A.9: **Tenure of Employees Who Left the Firm**

Notes: This figure displays the distribution of tenure length of employees who have left the company as of the end of the sample (2021), split between EWA users and non-users.

Panel A: Employees starting 3 months or less before EWA is offered



Panel B: Employees starting after EWA is offered

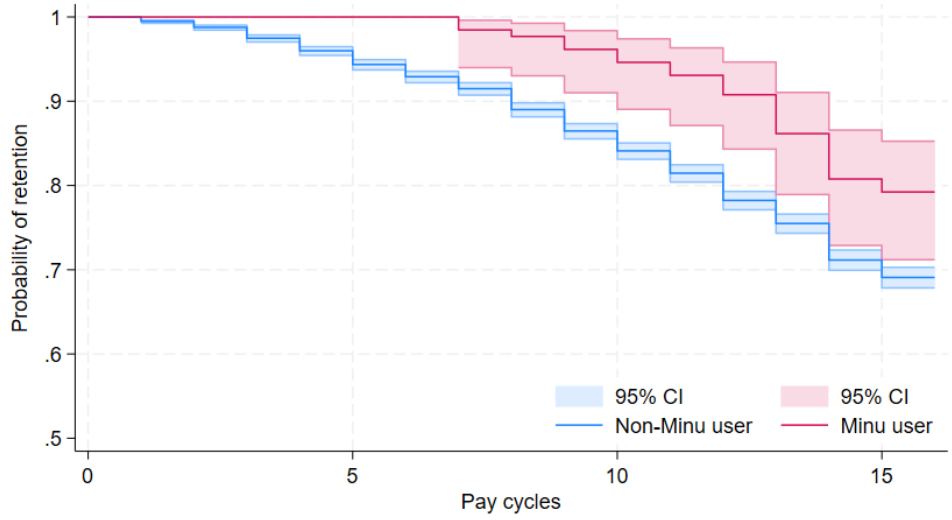


Figure A.10: **Turnover and EWA Use**

Notes: These figures display the share of employees who are still working at their initial employer in a given month since they joined, for EWA users and non-users. Panel A restricts the sample to the group of employees who started 3 months or less before EWA was initially introduced at the employer level, and panel B restricts to the employees who started working at the company after EWA is introduced.

A.4 Additional Tables

This appendix collects the robustness tables for the retention analysis of Section 6: Appendix Table A.1 re-estimates the extensive-margin hazard as a discrete-time model, and Appendix Table A.2 relates turnover to the usage pattern rather than to usage itself.

Table A.1: **EWA Use and Turnover: Discrete-Time Hazard Models**

	Leaves firm=1			
	(1) cloglog	(2) cloglog	(3) cloglog	(4) logit
Used Minu	0.779***	0.885***	0.897**	0.894**
	(0.036)	(0.042)	(0.043)	(0.044)
Male			1.080***	1.082***
			(0.029)	(0.030)
Age			0.995***	0.995***
			(0.002)	(0.002)
Baseline hazard	freq x hm	freq x hm	freq x hm	freq x hm
Start-quarter dummies	N	Y	Y	Y
Demographics+rank	N	N	Y	Y
Observations	411,859	408,251	249,502	249,502

Notes: This table reports a discrete-time counterpart of the extensive-margin hazard results of Table 5. The unit of observation is an employee–half-month row inside the risk set, treated as a Bernoulli trial for separation; the half-month grid carries heavy ties, which makes a grouped-time model a natural check on the Cox estimates. Columns 1–3 use a complementary log-log link, the exact grouped-time proportional-hazards analogue of the Cox model; column 4 uses a logit link, whose odds ratios approximate hazard ratios when failures per cell are rare. Reported figures are exponentiated coefficients, so a value below one indicates a lower separation hazard; this differs from Tables 5 and 6, which report log hazard ratios. The baseline hazard is a full set of company × pay-frequency × half-month dummies, the discrete-time counterpart of the Cox strata. Start quarter enters as dummies rather than interacted with the baseline, which is coarser than the corresponding Cox stratification. Standard errors, clustered on the employee, are in parentheses.

Table A.2: **Regular EWA Users and Turnover**

	Leaves firm=1		
	(1)	(2)	(3)
Regular (vs non-users)	-0.147	-0.190	-0.302**
	(0.150)	(0.152)	(0.151)
Other users (vs non-users)	0.040	0.025	0.006
	(0.084)	(0.084)	(0.083)
Male		0.064**	0.073***
		(0.026)	(0.027)
Age		-0.002	-0.004**
		(0.002)	(0.002)
Monthly salary			0.000***
			(0.000)
Company x Frequency Strata	Y	Y	Y
Start Quarter Strata	Y	Y	Y
Job Rank	N	strata	N
Gender	N	Y	Y
Number of Employees	42,641	26,239	21,757
Observations	404,375	245,148	205,792

Notes: This table presents coefficients (log hazard ratios) from Cox proportional hazard regressions relating the hazard of a worker leaving their firm to her EWA usage pattern rather than to usage itself; the implied proportional change in the hazard is $e^{\hat{\beta}} - 1$. Users are classified over the six pay cycles following their first withdrawal: *Regular* users are active in at least four of those six cycles, and *Other* users, the pooled remainder, are those who withdraw in three cycles or fewer. Group membership is time-varying and switches on only once the classification window closes, six pay cycles after first use, so that a user's earlier person-time contributes to the omitted category rather than to her eventual group; the omitted category is therefore non-users together with all pre-classification person-time. Users whose classification window is not fully observed are dropped. Because *Other* pools two behaviorally distinct groups, occasional users who withdraw against a large expense and users who try the service and stop, only the *Regular* contrast should be read as informative. Baseline hazards are stratified by company × pay frequency and by starting quarter, and additionally by job rank in column 2. Standard errors are clustered on the employee. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.