Protective Rejections in Consumer Credit

Alex P. Günsberg
Hanken School of Economics and
UNC Kenan-Flagler Business School (Visiting) *

First version: January 2025.
This version: May 10, 2026.
Online appendix: osf.io/mb3nf.
Author website: alexgunsberg.com.

Abstract

Some rejected borrowers look safe because no loan is disbursed. Credit-access diagnostics that count rejected applicants who later avoid default as missed opportunities observe default only after rejection, not default after disbursement. I estimate this counterfactual by linking Finnish comparison-platform applications to credit-registry defaults and postal-code data. Among applicants who already received lender offers, ordinary weekday processing frictions after the offer shift whether final disbursement occurs. On this margin, disbursement raises 18-month default by 55.6 percentage points ($p = 0.006$). These rejections are protective: disbursement itself creates the default they prevent. The economic stakes are large because debt contracts have capped upside and large downside exposure: one induced default erases the profits from roughly six fully repaid loans, with a five-to-ten-loan range across 7–11% cost-of-capital assumptions. Across area-income splits, the LMI estimate is 49.0 percentage points ($p = 0.006$, KP $F = 27.1$); in high-income areas, weekday weakly moves completion (KP $F = 4.5$) and the -7.7 percentage-point estimate is statistically indistinguishable from zero ($p = 0.812$). Standard approved-sample observables do not separate the harmful margin. Group gaps show where constraints are concentrated; they do not show whether rejected applicants would repay if funded. Policy needs that counterfactual. Removing local repayment signals can fund bad marginal loans, reject viable marginal applicants, and shrink the viable credit market when those signals predict repayment constraints.

Keywords: Consumer Credit, Credit Screening, Loan Default, Credit Access, Instrumental Variables

*Email: alex.gunsberg@hanken.fi. The author is currently visiting the UNC Kenan–Flagler Business School. ORCID: 0000-0002-2854-7628.

I. Introduction

Some rejected borrowers look safe because no loan is disbursed. Credit-access diagnostics often treat rejected applicants who later avoid default as missed credit opportunities. But that evidence observes default only after rejection, not default after disbursement. Credit is not just a test that reveals risk. It can change risk.

That distinction is the paper's object. A rejected applicant who avoids default has honored existing commitments without this loan. Calling the rejection false assumes she would also have avoided default with it. On the margin I estimate, that assumption fails: loan disbursement changes the borrower's balance sheet and repayment obligation enough to create default.

I estimate this counterfactual by linking Finnish comparison-platform applications to credit-registry defaults and postal-code data. The platform records the full funnel: applications, lender offers, borrower completion, and disbursement. The credit registry measures later default on any financial commitment for rejected applicants, approved applicants who do not complete, and approved applicants whose loans are disbursed. A lender mainly observes repayment on its own disbursed loans. The merged data make rejection, non-disbursement after approval, disbursement, and later default visible in one design.

The design uses ordinary weekday processing frictions after a lender has made an offer. Later-week applications move more slowly through final verification and completion. Among offered applicants, these frictions shift whether final disbursement occurs while leaving the lender's offer decision already made. The design therefore estimates the effect of disbursement for offered applicants whose completion responds to these last-stage frictions.

On this identified margin, disbursement raises 18-month default by 55.6 percentage points ($p = 0.006$). Rejection is protective here: it prevents defaults that disbursement itself would create. The label changes the policy action. A false-rejection rule would push more credit to these applicants; the causal estimate says that disbursement creates distress for them.

Those prevented defaults matter for portfolios. Debt contracts have capped upside and large downside losses. Using observed lender contract terms and an opportunity-cost benchmark, one induced default absorbs profits from roughly six fully repaid loans; across a 7–11% cost-of-capital range, the figure is five to ten repaid loans. Although the empirical laboratory is Finnish consumer credit, this arithmetic is not special to that market. Debt contracts generally expose lenders to capped upside and large downside default risk, so misclassifying protective non-disbursement as false rejection can distort credit supply beyond this setting.

The policy mistake is a wrong counterfactual. Outcome-based fairness and credit-access exercises define creditworthiness from observed or predicted repayment (Hardt, Price, and Srebro 2016; Kleinberg et al. 2018). In the closest lending exercise, Meursault et al. (2022) lower LMI-area approval thresholds to equalize true-positive rates, where the positive label is repayment. Their exercise also assumes that granting an additional loan does not change default probability. Fuster et al. (2022) show that machine learning improves mortgage-default prediction and changes predicted risk and pricing unevenly across race groups; their decomposition attributes much of the unequal incidence to greater flexibility in mapping permissible observables into default risk rather than to pure triangulation of race. These papers ask who is predicted to default and who receives credit. They do not estimate whether lending itself changes the default label used to classify false rejections. This paper estimates that missing treatment effect.

The same logic changes how screening evidence should be read. Credit decisions are made applicant by applicant. A lender has reason to lend to every applicant whose loan is expected to be viable, absent a binding portfolio constraint. Group gaps are descriptive. They show where constraints concentrate and can guide support, but they do not prove that a marginal rejected applicant would have repaid if funded. That requires the causal comparison estimated here.

The protective margin is visible where the instrument is strong. In LMI areas, disbursement raises default by 49.0 percentage points ($p = 0.006$; KP $F = 27.1$). In high-income areas, weekday weakly moves completion (KP $F = 4.5$) and the estimate is -7.7 percentage points ($p = 0.812$), so that split does not identify a comparable harmful effect. Weekday frictions also reduce completion most strongly in LMI areas. Postal codes are not randomly assigned, and they are not geography-as-destiny. They bundle sorting on resources, constraints, habits, networks, and fallback capacity that are not fully measured in the application file. Omitting this information does not remove local repayment risk; it worsens prediction and increases mistakes on both sides of the approval cutoff.

Within the already-approved offer sample, this is a demanding post-screen comparison. The lender has already screened these applicants, the effect does not rise with projected debt-to-income, and weekday-shifted borrowers have small standardized differences from the broader offer pool on credit score, age, income, and projected debt-to-income. Standard approved-sample observables therefore do not separate the borrowers for whom disbursement creates default.

The size of the viable market depends on the same margin. Better local repayment signals expand credit when they separate borrowers who would repay from borrowers whom disbursement would harm. Removing such signals creates two mistakes: bad marginal loans are disbursed, and viable marginal applicants are rejected. Both shrink viable credit. Defaults consume lender capacity; false rejections leave profitable loans unmade. Because market size is set at the approval margin, contraction and expansion are borne first by applicants closest to the threshold, including low- and moderate-income or protected-status-correlated groups that access policy is trying to help.

The registry-linked design brings rejected applicants, approved non-disbursed applicants, and disbursed borrowers into one empirical setting; replaces the false-rejection benchmark with the causal effect of disbursement; and shows that some rejections are protective because disbursement itself creates default on a margin moved by access policies and platform design.

The implication is not that less credit is always better. It is that policy needs the right counterfactual. Ex post non-default after rejection is the wrong false-rejection benchmark on margins where loan disbursement itself creates default. On the local disbursement margin identified here, policies that treat these rejections as mistakes target the wrong borrowers. The online appendix shows that, under transparent lender-response rules, pushing lending onto this IV-identified risky margin lowers portfolio value and creates incentives to tighten or reprice elsewhere.

The closest marginal-credit evidence studies the consequences of expanding consumer credit. Studies of payday and high-cost credit estimate that access increases household distress (Melzer 2011), that restricting access shifts borrowers toward inferior substitutes and worsens financial condition (Zinman 2010), and that present bias and naivete distort payday borrowing and shape the welfare effects of regulation (Allcott et al. 2022). Closest in setting, Dobbie et al. (2021) show that marginal take-up in U.K. high-cost lending raises default by 44.7–58.8 percentage points. This paper differs by isolating disbursement among applicants who have already received an offer and by linking platform records to credit-registry outcomes for rejected applicants, approved non-takers, and approved takers, so it tests whether rejection on this margin is protective rather than false.

The paper also connects screening, algorithmic lending, and information constraints. Classic work shows how imperfect information rations credit (Stiglitz and Weiss 1981). In consumer-credit markets, empirical work documents adverse selection, liquidity constraints, and moral hazard (Adams, Einav, and Levin 2009; Karlan and Zinman 2009; DeFusco, Tang, and Yannelis 2022), and financial innovation expands credit to riskier borrowers by enabling new contracts targeted at marginal households (Livshits, Mac Gee, and Tertilt 2016). Recent work shows the upside of richer information: digital footprints predict consumer default and affect access to credit (Berg et al. 2020), machine-learning models expand lending at a given delinquency rate (Tantri 2021), and automated fintech credit relaxes information frictions for platform vendors (Hau et al. 2024). Other work shows how machine learning changes mortgage-credit allocation across groups (Fuster et al. 2022) or reveals pricing disparities in lending (Bartlett et al. 2022).

This paper studies the policy risk on the other side of that literature. Richer variables expand credit when they improve screening. Restrictions or equalization of repayment-relevant variables remove information that protects borrowers and portfolios when those variables predict repayment constraints. The mechanism is worse prediction: lenders disburse more bad marginal loans and reject more viable marginal applicants. Both errors shrink the viable market. Better full-funnel training data work in the opposite direction: they distinguish viable borrowers inside groups where constraints are concentrated, preserve market size, and create room to lend closer to the margin. The object here is therefore not prediction or group gaps alone, but whether observed rejections at the margin are mistakes or instead prevent defaults that loan disbursement itself would have caused.

II. Institutional Setting

The setting is the Finnish high-cost unsecured consumer-credit segment. These are not collateralized consumer loans priced like bank credit. They are large unsecured obligations with average APR 17.3%, modal APR 21.8%, average maturity 75 months, and loan sizes up to EUR 60,000. That combination matters for the paper. In a product this expensive and long-lived, the difference between loan disbursement and non-disbursement creates scope for large downstream effects on distress, and the lender's payoff asymmetry is correspondingly steep.

The data come from a Finnish online comparison platform. Applicants actively seek credit, submit one application, receive offers from multiple lenders, and then decide whether to complete the transaction. They are not passive borrowers being funneled into one product. They are shopping, comparing, and deciding whether to follow through after the initial offer stage. That structure isolates the final completion decision among applicants who are already searching for high-cost unsecured credit.

This market combines high-cost unsecured debt with a digital comparison-platform setting, richer linked data, and formal affordability checks. That combination is precisely why the setting is valuable: it studies a dangerous credit product while observing later default outcomes for the full applicant pool, not only funded borrowers.

The platform also separates several decisions bundled together in lender data. First, the borrower decides to search for credit. Second, lenders decide whether to extend offers. Third, approved borrowers decide whether to complete the transaction once the exact timing of verification and disbursement becomes concrete. The weekday design operates on that last step. This separation matters because many product-design and policy interventions work through completion convenience at the end of the borrowing process.

Weekday frictions arise naturally in this setting. Even digitally originated loans still pass through approval queues, anti-fraud checks, identity confirmation, and payout processes that are slower later in the week. Friday applications face weekend payout delays that Monday applications avoid. These delays are not artificial researcher-generated frictions. They are ordinary operational constraints that change how quickly applicants can finish the transaction.

A. The Cost Structure of Unsecured Lending

The cost structure of unsecured lending helps explain why this margin matters so much. A successful loan generates profits gradually over time, whereas a bad loan produces a large loss because there is no collateral to cushion default. Table 1 quantifies this asymmetry using the average loan in the sample: amount EUR 10,626, term 75 months, and APR 17.3%. Assuming a lender cost of capital of 8.0% and operating costs of EUR 531, the discounted payoff from a fully repaid average loan is EUR 2,531. A default at the average default timing of month 13 implies an origination-date default alternative cost of roughly EUR 13,828. One default therefore absorbs the profit from roughly 6 fully repaid loans. Over the 7.0–11.0% cost-of-capital range in Panel B, the ratio is 5–10 fully repaid loans.

Table 1: Default Alternative Cost Benchmark

Panel A: Average Loan and Max Loan Scenarios			
		Max Loan Amount	
	Average Loan	Immediate Default	Avg Default
Lender Cost of Capital (%)	8.0	8.0	8.0
Alternative Lending Return (%)	17.3	17.3	17.3
Loan Amount (EUR)	10,626	60,000	60,000
Months to Default	13	0	13
Unpaid Principal at Default (EUR)	9,497	60,000	53,620
Default Alternative Cost at Default (EUR)	15,076	99,444	85,122
Default Alternative Cost at Origination (EUR)	13,828	99,444	78,078
Discounted Payoff, Repaid Avg Loan (EUR)	2,531	2,531	2,531
Repaid Loans Needed	6	40	31

Panel B: Average Loan Sensitivity to Cost of Capital					
Lender Cost of Capital (%)	7.0	8.0	9.0	10.0	11.0
Default Alternative Cost at Origination (EUR)	13,978	13,828	13,680	13,534	13,389
Discounted Payoff, Repaid Avg Loan (EUR)	2,934	2,531	2,143	1,771	1,413
Repaid Loans Needed	5	6	7	8	10

Notes: The average loan has an amount of EUR 10,626, a term of 75 months, and a mean APR of 17.3% (mode 21.8%); for defaulting loans, the average time to default is 13 months (Panel A, column 1). Maximum loan amount is EUR 60,000. The borrower is assumed to make scheduled payments through the listed default month, so no pre-default payments are treated as lost. 'Default Alternative Cost at Default' is the unpaid balance at default plus the present value, measured at the default date, of the alternative lending return that the unpaid capital could have generated over the remaining contractual term. The alternative lending return is the average APR; it is not the lender's cost of capital. 'Default Alternative Cost at Origination' discounts that default-date cost back to origination using the lender's cost of capital. 'Repaid Loans Needed' divides this origination-date opportunity-cost benchmark by the discounted payoff from a fully repaid average loan, after operating costs (EUR 531, 5% of the average loan amount), and rounds up. This is an opportunity-cost benchmark, not an accounting charge-off. Panel B varies the lender cost of capital over a fintech-relevant range; across 7–11%, one default absorbs the profit from 5–10 fully repaid average loans. Assumes zero recovery of unpaid principal and homogeneity in loan characteristics.

This asymmetry matters for more than lender accounting. Harmful marginal lending is expensive enough that lenders have incentives to protect portfolio value if a policy or platform design raises completion for borrowers harmed by disbursement. The appendix uses an illustrative lender-response exercise to quantify that downstream incentive logic under stated assumptions.

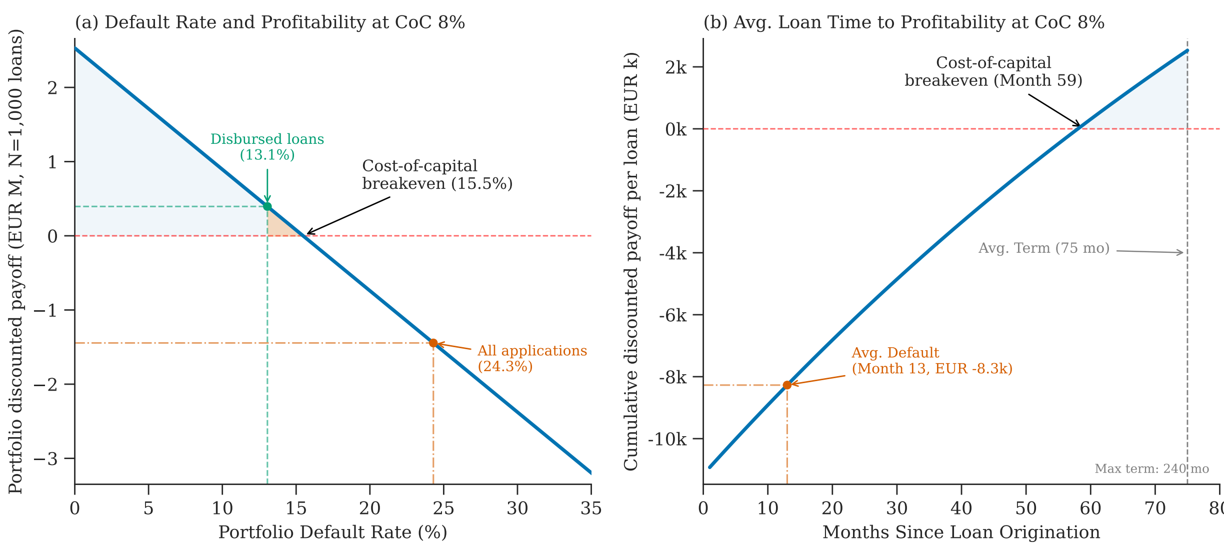

Figure 1: Lender Profitability and Breakeven Analysis

Notes: This figure illustrates the sensitivity of lender profitability to default risk from both a portfolio and a single-loan perspective. The analysis is based on the average loan from the sample (amount: EUR 10,626; term: 75 months; APR: 17.3%) and assumes a 8.0% lender cost of capital. **Panel (a)** plots expected net profit for a portfolio of loans as a function of the default rate. The figure marks the breakeven default rate of 15.5% and the observed default rate of disbursed loans, 13.1%. **Panel (b)** shows the cumulative discounted payoff for a repaid average loan over time and marks the breakeven month, 59.

III. Interpretive Framework

Two facts matter for interpreting the estimates.

First, rejection is the operative screening choice when debt payoffs are asymmetric. A higher interest rate cannot fully undo default risk when the lender's upside is capped and default losses are large. Let X_i denote the information the lender observes at application and let $\hat{p}_i = f(X_i)$ denote the lender's predicted default risk based on those observables. For a loan with terms L_i, expected profit is

$$\Pi_i(r_i) = (1 - \hat{p}_i)R(L_i, r_i) - \hat{p}_i M(L_i, r_i),$$

where $R(L_i, r_i)$ is the value of a repaid loan and $M(L_i, r_i)$ is the loss from default. This is a general debt-contract problem, not a platform-specific oddity: contractual upside is capped, while default exposes the lender to principal loss, recovery costs, and lost interest. In this unsecured-credit setting, those default losses are large relative to repayment profits. That is why the lender's relevant choice is a discrete approval threshold,

$$\text{Approve}_i^{\text{priv}} = \mathbf{1}\{\Pi_i(r_i) \geq 0\},$$

as a discrete screening choice. Even when competitive pricing reduces excess expected profits and moves risk-adjusted returns toward equality, realized payoffs remain asymmetric. Higher interest rates also raise payment burdens and make repayment less likely. In that environment, rejection screens out some applicants more effectively than a higher price.

The decision rule is individual. It does not require rationing one applicant because another applicant receives credit. Absent a binding capital or portfolio constraint, every applicant whose expected payoff clears the break-even condition receives a loan. Repayment-relevant information is therefore market-expanding when it separates viable borrowers from borrowers whose expected downside exceeds capped upside. Group averages are not the decision rule. They describe the incidence of individual risk scores across coarse categories.

One reason this threshold representation is operationally natural is that repayment risk is learned from binary realized outcomes. A lender assigns a continuous default probability to an applicant, but each originated loan supplies only a default or non-default draw, and lender-specific learning comes from aggregating those binomial realizations. The lender also observes those draws only for loans it books. In a high-dimensional applicant space, especially on a platform where lenders extend offers but do not originate every offered loan, many lender-specific regions of the risk distribution are thin. Individualized prices therefore cannot be treated as a frictionless mapping from every covariate profile to its exact default risk. Cutoffs are a practical way to act on noisy granular risk estimates when default losses are large and higher prices raise repayment burdens.

The platform is therefore a useful laboratory for a general screening problem. Unsecured consumer loans generate many small, near-term downside realizations. Other debt markets, such as mortgage credit, generate rarer but much larger and more correlated tail losses. The empirical object differs, but the screening problem is the same: when downside losses are asymmetric and borrower-level risk is learned noisily from realized defaults, rejection is an efficient screening response.

Second, ex post non-default among rejected applicants is the wrong benchmark for evaluating rejection quality once loan disbursement changes repayment outcomes. Let $Y_i(1)$ denote the applicant's later default outcome if the loan is disbursed and $Y_i(0)$ the outcome if it is not. Observing that a rejected applicant later remains current reveals something about $Y_i(0)$ for people whose loan is not disbursed. It does not reveal whether disbursement would have left that same applicant equally safe. The paper's design identifies exactly that missing counterfactual for the borrowers whose completion decision changes with weekday processing frictions: whether disbursement would have changed their later repayment path.

IV. Data and Research Design

A. Data

The dataset contains 758,622 loan applications from 214,539 individuals between 5 April 2019 and 15 June 2024. Across these applications, 39 lenders make 13,033,771 credit decisions and extend 783,983 offers, for an overall offer rate of 6.0% (Table 3). The central empirical advantage comes from merging platform records with credit-registry outcomes. The registry measures default on any financial commitment for rejected applicants, approved non-takers, and approved takers. Ordinary lenders do not observe this object. A lender mainly observes repayment outcomes on loans it actually disburses; it does not observe the later outcomes of every applicant rejected by other lenders or every approved borrower who walks away before disbursement. Because the merged data cover the full applicant pool, the paper tests whether disbursement itself changes repayment. The observable set is also unusually rich relative to ordinary lender data. It combines applicant disclosures, credit bureau records, postal-code and municipal registry joins, the full record of cross-lender decisions on each application, and application-funnel behavior. Section VI returns to this point: the "what observables miss" claim is a claim about an observable set richer than what single lenders ordinarily see.

Table 2: Three Levels of Inference in Rejection Audits

Level	Data observed	Inference
Level 1: typical lender	Own disbursed loans	Default among funded borrowers only
Level 2: registry merge	Default for rejected applicants, non-takers, and takers	Outcome frame for the full applicant pool
Level 3: registry + IV	Exogenous shift in disbursement after offer	$E[Y_i(1) - Y_i(0) \mid C]$, the disbursement effect for the weekday-moved offered-applicant margin

Notes: C denotes offered applicants whose completion responds to weekday processing frictions. Level 3 does not identify $Y_i(1)$ for rejected applicants generally; it identifies the local disbursement effect on the weekday-moved offered-applicant margin.

The sample represents 4.7% of Finland's adult population. Applicants have median annual income of EUR 24,600 versus EUR 25,398 nationally, median age 40.0 versus 52.8 nationally, and lower average debt than the national average (EUR 16,551 versus EUR 21,443). They are also somewhat less likely to hold higher-education degrees than the national population (Table 4). Following Meursault et al. (2022), I classify postal-code areas as low- and moderate-income if median area income is at or below 80% of the national median; the remaining areas are classified as higher-income. This split is used later to show where the harmful disbursement margin is concentrated.

Table 3: Summary Statistics for Loan Platform Activity: 5 April 2019–15 June 2024

Statistic	**Count**	**Percentage**
Applicants		
Unique Applicants	214,539	—
With Any Default	74,116	34.5%
Applications and Decisions		
Loan Applications Submitted	758,622	—
Decisions per Application	∼17	—
Lender Decisions	13,033,771	—
Offers Extended	783,983	6.0%
Loans		
Funded Loans	46,161	6.1%
In Default	5,995	13.0%

Notes: Summary statistics from the raw dataset, requiring valid application and applicant identifiers and submission timestamps. Percentages: defaults among unique applicants; offers among lender decisions; funded loans among applications; defaults among funded loans. 'Any Default' includes any payment default recorded for an individual regardless of loan approval. Average decisions per application ≈17.2 (13,033,771 / 758,622). There are 39 unique lenders.

Table 4: Sample Demographics versus Finnish Population

Characteristic	Sample	Finland (2023)
Undergraduates (%)	15.9	22.4
Graduates (%)	5.7	12.9
Median Age	40.0	52.8
Median Annual Income (EUR)	24,600	25,398
Average Debt (EUR)	16,551	21,443

Notes: 214,539 unique individuals, representing approximately 4.7% of Finland's adult population (aged 18+ in 2023). Education percentages are for individuals aged 18 or older. Median income is annual after-tax (sample: 12× monthly net; Finland: 2023). Average debt includes unsecured loans, vehicle loans, and credit card debt. Finnish national figures: Statistics Finland (2023).

B. Outcome Measures

For the causal analysis, the outcome is default on any financial commitment within 540 days (18 months) of the application. This horizon is a design choice, not a caveat: weekday changes processing speed and disbursement, disbursement changes the borrower's debt path, and financial distress takes time to appear in recorded defaults. For the policy application, I use the broader default indicator observed in the merged credit-registry data for all applicants.

Table 5 summarizes approved loans. Default timing in these loans is informative for the break-even calculations reported in Section II.A, but it is also measured in a sample subject to pandemic-era forbearance and right-tail truncation near the end of the panel. Both forces make the observed timing of default conservative for the purpose of measuring lender losses.

Table 5: Annual Characteristics of Approved Loan Offers

	2019†	2020	2021	2022	2023	2024†
Offered Amount Avg (EUR)	6,590	8,293	10,483	10,874	11,210	10,977
Offered Term Avg (months)	68	69	72	77	79	76
Months to Default, 50th pct	19	21	18	13	7	2
Months to Default, 90th pct	48	41	31	23	12	4

Notes: Statistics for lender loan offers, including those not funded. Default timing reflects disbursed loans only. Dashes indicate no defaults. † Partial years (8.9 and 6.5 months).

The ability to observe defaults for rejected as well as approved applicants is what makes the paper possible. The credit-registry merge measures later default for the full applicant pool, including applicants who never book a loan with a given lender. In ordinary lender records, credit audits treat a rejected applicant who later remains current as evidence that the lender made a mistake. That interpretation assumes that taking the loan would have left the applicant's later repayment path unchanged. The design tests that assumption on the weekday-moved offered-applicant disbursement margin: when disbursement changes after an offer, does later default risk change too?

C. Research Design: Weekday Frictions as Timing Variation

The causal question is whether loan disbursement changes later default. People who complete the process differ from people who do not in ways the data do not fully capture, so the paper uses application weekday to shift completion through processing speed.

Application → offer → completion → disbursement → default

Weekday processing frictions move completion and disbursement after offer.

Figure 2: Application Funnel and the Weekday-Moved Margin

Applications submitted later in the week move more slowly through approval, final verification, and disbursement. That extra time slightly lowers completion. Some offered applicants who would have taken the loan if cash had arrived quickly do not complete the transaction when the process slows down. Others proceed regardless. The paper uses that difference to estimate the causal effect of loan disbursement for the marginal group whose behavior changes with those frictions.

The statistical approach has two steps. In the first step, I predict whether the loan is completed using the day of the week the application arrived, together with application and platform controls such as credit score and income. In the second step, I use only the part of completion explained by the day of the week to estimate the effect of loan disbursement on later default. This isolates the effect for applicants whose behavior changes with small processing delays. For comparison, the simple comparison that does not isolate this weekday-driven component replaces the predicted completion with realized completion. Controls include rich application and platform observables: credit score, projected debt-to-income, expenses relative to a reference budget, loan characteristics, family structure, applicant history, measures of stability, and year-quarter fixed effects. Continuous controls are standardized using the estimation-sample mean and standard deviation before entering the IV specifications. This scaling leaves the treatment estimand unchanged but puts continuous covariates on a common scale and makes the projected-DTI interaction interpretable as a one-standard-deviation change. Treatment, outcome, weekday instruments, fixed effects, and binary indicators are kept in their original units. Standard errors are clustered at the borrower level.

Formally, the first step relates completion/disbursement to weekday and controls,

$$D_i = \alpha + Z_i'\pi + X_i'\gamma + u_i,$$

where D_i is loan disbursement after completion, Z_i is the vector of weekday indicators, and X_i is the control set. The second step relates default to the part of completion/disbursement explained by weekday,

$$Y_i = \beta \widehat{D}_i + X_i'\delta + \varepsilon_i,$$

where Y_i is default within 540 days. In plain English, the first equation predicts whether the loan is completed using the day of application and application and platform observables; the second uses only the part of completion explained by weekday timing to estimate the effect of loan disbursement on later default.

The instrument Z_i is a vector of weekday dummies, with Monday omitted. The estimand is the causal effect of loan disbursement for borrowers whose completion decision changes when weekday frictions change. In potential-outcomes notation,

$$\beta = E\left[Y_i(1) - Y_i(0) \mid D_i(\text{fast}) > D_i(\text{slow})\right],$$

where $Y_i(1)$ and $Y_i(0)$ are the borrower's potential default outcomes with and without loan disbursement and $D_i(\text{fast})$ and $D_i(\text{slow})$ are the corresponding completion choices under faster and slower processing. The identifying assumption is that slower processing does not make any borrower more likely to complete the loan. In plain English, the coefficient applies to offered applicants who take the loan when the process is relatively easy but not when it is relatively slow.

The sample for the causal analysis includes applications submitted by June 2023 so that 18-

month default outcomes are fully observed. Observed default rates already suggest little raw selection between takers and non-takers: the default rate is 12.04% in the full sample, 11.48% among non-takers, and 17.7% among takers (Table 6). The weekday design then isolates the causal component of completion and disbursement induced by processing frictions.

Table 6: Default Rates by Loan Take-Up Status

Group	Default Rate (%)	N
All offered applicants	12.04	92,198
Non-takers (offered, did not take up)	11.48	83,993
Takers (offered and took up)	17.7	8,205

Notes: Default within 18 months (540 days) on any bill, conditional on receiving a loan offer. Sample restricted to applications submitted at least 540 days before the last observation date to ensure complete outcome windows.

C1. Why Weekday Is a Plausible Source of Variation

A credible design needs two ingredients: weekday predicts loan disbursement, and weekday does not independently sort applicants by untreated default risk. The relevant challenge is not that the affected margin is selected—the LATE is defined for applicants whose completion responds to timing. The challenge is whether Monday–Thursday and Friday–Sunday applicants would have defaulted differently even without disbursement, including because fraudulent or verification-failing applicants arrive on particular days. The diagnostics below target that threat.

Completion changes over the week. Completion is high early in the week and lower on later-week, higher-friction days: it is 22% on Monday and 21% on Saturday, a gap of 1.8 percentage points, or 8% relative to Monday (Figure 3). The timing lines up with the operational channel. Relative to Monday, Friday applications take +22.35 more hours to reach payout and Saturday applications +26.93 more hours; in the post-selection window, Friday and Saturday still add +18.84 and +18.22 hours, respectively. The online appendix reports the delay table. Completion is lower on the same higher-friction days where payout takes longer.

The balance evidence supports the timing interpretation. The applicant-composition critique predicts weekday differences in pre-determined risk or default absent disbursement. Four facts point instead to processing timing. First, the reduced form for 18-month default tracks the first step: later-week days reduce completion and also predict lower later default. Second, weekday has no predictive power at horizons too short for loan-induced distress to materialize; those short-horizon nulls are placebo-style support for the design. Third, pre-determined observables change little across weekdays. The online appendix reports the balance table comparing credit score, age, income, debt, residence tenure, occupation tenure, underage children, and first-application status across weekdays. The largest scaled difference in means across any weekday–Monday pair is 0.0656, below the 0.10 imbalance benchmark. Fourth, those weekday differences are small relative to the differences associated with realized loan disbursement. The median weekday imbalance is only 42.3% of the corresponding taker-versus-non-taker imbalance.

Finnish payroll timing also makes the most obvious payday channel unlikely. Pay schedules are not tied to a fixed weekday pattern, and banking-day adjustments further rotate paydays across weekdays. Application weekday is therefore not systematically tied to paycheck timing or payday liquidity shocks. The weekday channel is processing delay. Mean distance to the nearest adjusted payday differs little across weekdays, and controlling directly for payday proximity leaves the main estimate at 55.8 percentage points; the online appendix reports the payday balance and robustness tables.

A separate concern is weekend fraud or failed verification. Because the sample is already post-offer, this story must be more specific than bad applicants applying on weekends: Friday–Sunday offered applicants would need to have higher default risk even if no loan is disbursed. The non-disbursed cells go the other way. Among non-disbursed applications, Friday–Sunday default is 12.10%, compared with 12.51% for Monday–Thursday. The Kitagawa-style test for the coarse Monday–Thursday versus Friday–Sunday instrument has one point-estimate sign violation, in the defaulting non-disbursed cell; the maximum violation is small (0.00183), and the bootstrap test does not reject ($p = 0.497$). If real, this is the direction that would make the IV estimate too large, so it is the right cell to examine. The point estimate is statistically indistinguishable from zero and goes opposite the simple weekend-fraud story.

The overidentification test in the main specification is also not rejected (Hansen $p = 0.191$). The delay, short-horizon, balance, non-disbursed-default, payday, overidentification, and testable-inequality evidence give the same answer: weekday moves completion, not a visibly different borrower pool.

Design checks supporting the timing interpretation.

1. **Operational delay:** Friday and Saturday applications face payout delays relative to Monday.
2. **Completion:** completion is lower on later-week, higher-friction days relative to early-week days.
3. **Balance:** the maximum weekday SMD is below 0.10, and weekday imbalance is much smaller than taker–non-taker imbalance.
4. **Short-horizon timing:** weekday has no stable effect before loan-induced distress should appear.
5. **Payday robustness:** payday proximity does not move the result.
6. **Non-disbursed default:** Friday–Sunday non-disbursed applicants do not default more than Monday–Thursday non-disbursed applicants.

7. **Overidentification and testable inequalities:** the Hansen test is not rejected, and the Kitagawa-style diagnostic is not rejected.

8. **Same-borrower sign check:** the within-applicant estimate is positive but less precise.

A stricter design compares the same repeat applicant across different weekdays. That specification absorbs all time-invariant borrower traits and estimates a 42.5 percentage-point effect ($p = 0.101$; KP $F = 9.2$). The main design carries the evidentiary weight; the same-person specification is a positive but less precise check. The online appendix reports the full weak-instrument diagnostics.

Figure 3: Loan Completion by Application Weekday

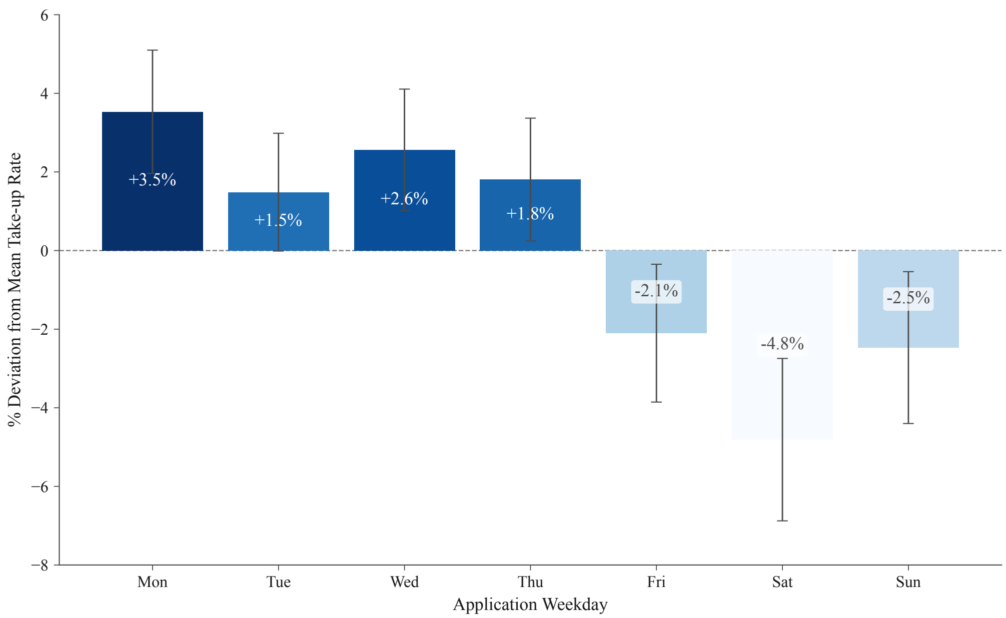

Notes: Mean loan completion rate by weekday of application. Completion declines from 22% on Monday to 21% on Saturday, a decline of 1.8 percentage points (8%). Bars show 95% confidence intervals. Sample size: 91,170 observations.

D. What the Design Identifies

The design identifies a specific causal effect for a specific group.

Who is identified. The borrowers whose behavior the design moves are offered applicants whose completion changes when weekday processing makes the final step a bit easier or a bit harder. They are active searchers who reach the end of the application pipeline and are close to indifferent between taking and not taking the loan. These are the borrowers whose disbursement status is shifted by processing speed.

What counterfactual is identified. The design recovers the effect of loan disbursement versus non-disbursement for offered applicants whose completion responds to weekday processing frictions. It is therefore the effect of loan disbursement for that group, not the effect of an offer abstractly defined apart from completion.

Why this is informative about rejection policy. Policy discussions treat non-default among rejected applicants as evidence of mistaken rejection. That logic assumes that taking the loan would not itself have changed repayment outcomes. The paper's design directly estimates that counterfactual for the relevant margin. On the identified disbursement margin among offered applicants whose completion responds to weekday processing frictions, loan disbursement raises 18-month default probability by 55.6 percentage points. Non-default without a loan is therefore not evidence of screening error on this margin. Calling those rejections false rejections misclassifies a protective causal effect as a mistake.

Scope. The result applies to the operational disbursement margin among offered applicants whose completion responds to weekday processing frictions. It identifies one operational channel through which rejection protects borrowers from taking loans that would worsen their repayment outcomes.

V. Main Results

A. Loan Disbursement Raises Default on the Identified Margin

Table 7 shows the causal effect of disbursement on the weekday-moved offered-applicant margin. Column 1 gives the baseline 2SLS estimate; column 2 adds the endogenous interaction between disbursement and projected debt-to-income; columns 3 and 4 split the sample by area income. Loan disbursement raises 18-month default probability by 55.6 percentage points. The first stage is strong enough for the main causal estimate (KP F = 22.03). For offered applicants whose completion falls when processing slows, disbursement creates about 55.6 extra defaults per 100 marginal disbursements.

Table 7: Loan Disbursement on the Weekday-Shifted Margin Increases Default Risk

	(1) Full sample	(2) + DTI interact.	(3) LMI areas	(4) High-income
Panel A: Treatment effect				
Loan disbursement	0.556***	0.573***	0.490***	-0.077
	(0.202)	(0.188)	(0.178)	(0.323)
Disbursement × Proj. DTI		0.021		
		(0.277)		
Panel B: Diagnostics				
First-stage F-statistic	22.0	38.6 / 13.6	27.1	4.5
Hansen J (p-value)	0.191	0.338	0.020	0.677
LIML coefficient	0.783			
Controls (44)	Yes	Yes	Yes	Yes
Quarter FE (15)	Yes	Yes	Yes	Yes
Observations	91,170	91,170	72,942	18,053
Individuals	55,975	55,975	45,178	11,466

Notes: Dependent variable: default on any bill within 540 days. Treatment is loan disbursement after an offer. All columns: 2SLS with weekday instruments (Tue–Sun, Monday omitted). Standard errors clustered by borrower in parentheses. Continuous controls are standardized using estimation-sample moments; treatment, outcome, weekday instruments, fixed effects, and binary indicators remain in their original units. Column (2) instruments both disbursement and disbursement × projected DTI using weekday dummies and their DTI interactions (12 instruments), so the interaction is interpreted per one-standard-deviation increase in projected DTI. Columns (3)–(4) split by postal-code median net income; LMI = bottom 80% (cutoff: EUR 26,434). Column (2) reports first-stage partial F-statistics for D_i / $D_i \times \widetilde{DTI}_i$, computed under borrower-clustered covariance; the smaller statistic is the binding relevance constraint. Panel B reports the LIML coefficient for the full-sample specification (41% larger than 2SLS). The online appendix reports the full specification with all control variables. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The credit-registry merge measures later defaults for all applicants, not just borrowers whose loans were disbursed. Combined with the weekday design, this lets the paper estimate what loan disbursement itself does to repayment for marginal offered borrowers. The borrower-level fact is direct: on this margin, a rejection that looks false under ex post non-default is causally protective once the disbursement counterfactual is estimated.

The estimate is local because credit policy and platform design operate at the margin. The design isolates offered applicants whose disbursement status is shifted by weekday processing frictions. It is the effect of loan disbursement for a specific operational margin: borrowers moved by verification, payout timing, reminder flows, or other last-stage frictions. Those are exactly the margins that platform design changes and many access-expansion interventions move in practice.

The magnitude is best read with the no-loan benchmark for the same margin. A near-threshold matching exercise approximates that benchmark from a proxy construction: using an XGBoost model trained on 132 observables to predict completion, non-takers whose predicted completion probability falls at or above the 75th percentile have an untreated 18-month default rate of 11.3% (95% CI 10.8–11.9; threshold sensitivity 7.1–11.8%; reported in the online appendix). Combining that proxy estimate of $E[Y(0) \mid C]$ with the IV estimate implies a with-loan default rate near 66.9%. Without the loan, this margin has a low untreated default rate: the near-threshold proxy puts $E[Y(0) \mid C]$ at 11.3%. With the loan, the implied default rate is near 66.9%. The loan is the difference, and the difference is not visible in the underwriting process that already approved these applicants. In plain units, 100 extra loan disbursements on this margin imply about 55.6 extra defaults. Using observed lender-set contract terms and the opportunity-cost benchmark in Table 1, each default absorbs the profit from roughly 6 fully repaid loans at the baseline 8.0% cost of capital, and 5–10 loans over the 7.0–11.0% range. The estimate is economically large even though the complier group is local. The matching exercise anchors the magnitude; the weekday IV remains the source of the causal treatment effect.

The online appendix reports an alternative limited-information estimator that is larger than the baseline estimate. The LIML estimate, 0.783, exceeds the baseline 2SLS estimate, 0.556. Since the conventional many-instrument concern would pull 2SLS toward OLS—here much smaller—the LIML comparison points away from upward many-IV inflation of the headline estimate. I use the baseline as the organizing estimate because it is the conservative specification; the limited-information result reinforces the sign and economic scale.

A simple way to interpret the result is to compare two offered applicants who search on the same platform and differ in processing speed because of application weekday. The Monday applicant receives cash quickly and completes the loan. The later-week applicant faces a longer wait and, at the margin, does not proceed to disbursement. The 18-month difference in later default between those two states is the object the design identifies: the real consequence of one additional unsecured loan reaching a marginal borrower's bank account.

The simple comparison that does not isolate the weekday-driven component is much smaller than the causal estimate. That is intuitive. Realized completion mixes safer borrowers who complete regardless with the more fragile group identified by weekday processing frictions. Once the design isolates the borrowers whose behavior actually changes at the margin, the effect is much larger. Raw repayment comparisons are therefore misleading for policy because they combine inframarginal borrowers and the risky disbursement margin as if they were the same object.

Figure 4 shows how the effect evolves over time. The dynamic profile is a design feature.

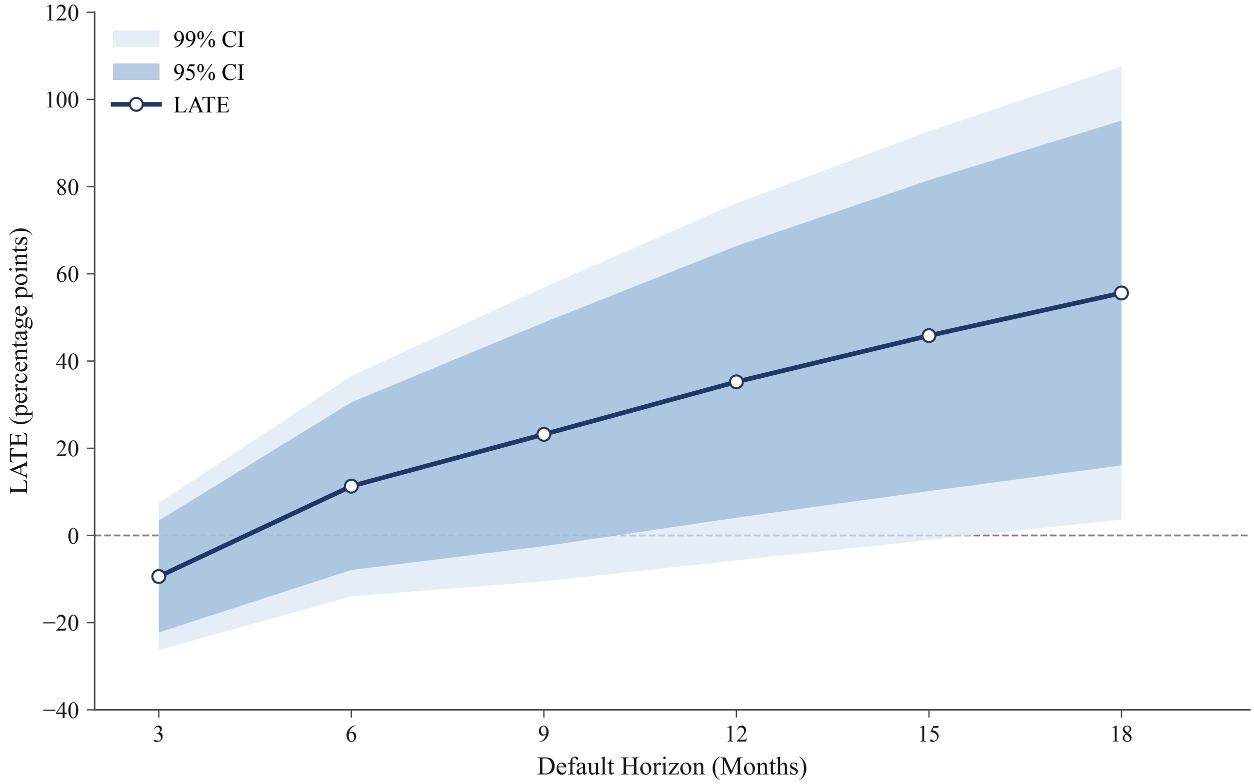

Figure 4: Treatment Effect of Loan Disbursement on Default, by Horizon
Notes: Each point is an estimate from the instrumental-variables baseline specification, varying only the outcome horizon. Shaded regions show 95% confidence intervals. Standard errors are clustered at the borrower level.

Weekday timing shifts processing and disbursement; disbursement changes the borrower's debt path; and financial distress takes time to appear in recorded default. The estimate has no stable predictive content at horizons too short for loan-induced distress to materialize, then builds steadily, crosses into clear net harm around months 10–11, and reaches full magnitude by 18 months. That path is more consistent with debt-related harm accumulating over time than with a bookkeeping artifact, a timing quirk in recording defaults, or a purely short-run liquidity story.

A stricter same-borrower specification in the appendix also stays positive. The repeat-applicant sample has a weaker relationship between day of week and completion, so the specification is a directional sign check on the main design.

The result is also externally plausible. Dobbie et al. (2021) provide the closest scale comparison: their examiner-leniency design studies the effect of approval on marginal high-cost borrowers. This paper studies a different margin—disbursement after an offer, shifted by operational weekday frictions. The scale comparison still matters because both designs show large causal lending effects, but the estimand is not the same. The present design isolates the disbursement-after-offer counterfactual that false-rejection audits miss.

VI. What Pre-Disbursement Variables Reveal

The sample is offered applicants. These borrowers have already cleared lender approval on credit score, income, projected debt burden, employment, residence, and the rest of the standard underwriting set. The tests ask what observed borrower variables and local-area information reveal about the harmful disbursement margin among borrowers commercial screening already approved.

LMI status locates the estimated harmful margin. In the area-income split, the LMI estimate is 49.0 percentage points ($p = 0.006$; KP $F = 27.1$). The high-income split has a weak first stage (KP $F = 4.5$) and a statistically insignificant -7.7 percentage-point estimate ($p = 0.812$). Within the already-approved offer sample, the main balance and affordability tests show that credit score, age, income, and projected debt-to-income do not separate the residual borrowers moved by weekday frictions. The machine-learning exercise has a different role: it predicts completion to construct a no-loan baseline for likely compliers; it does not evaluate default screening.

A. Measured Affordability Does Not Explain the Harmful Margin

Column 2 of Table 7 asks whether the effect is larger for borrowers with higher measured debt burden. The model regresses 18-month default on loan disbursement, loan disbursement interacted with standardized projected debt-to-income, projected debt-to-income itself, the baseline controls, and application-quarter fixed effects. Both endogenous terms are instrumented: disbursement is instrumented with weekday dummies, and the disbursement$\times$DTI interaction is instrumented with weekday dummies interacted with standardized projected debt-to-income. Projected debt-to-income is the applicant's post-loan debt burden measured before disbursement: loan amount plus existing debt minus the consolidation amount, divided by gross monthly income.

The interaction is the main affordability test because it keeps debt burden continuous and asks directly whether the treatment effect rises with measured repayment pressure. It does not. Both the projected-DTI gradient in the 2SLS specification reported in the online appendix and the disbursement–DTI interaction, 0.021 with $p = 0.939$, are near zero in the offered-applicant sample. The first fact is consistent with lenders having already screened out the projected-DTI variation that would predict default at the offer stage. The second says the disbursement harm is orthogonal to the residual DTI variation that remains among approved applicants. The harmful margin sits where DTI-based screening cannot reach, not within the DTI variation that survives approval. The harmful margin is therefore not picked out by projected DTI. Appendix checks that split the sample by leverage, exclude mortgage holders, and recompute projected debt-to-income net of mortgage exposure are robustness checks on the same affordability channel, not separate evidence that selects the preferred result.

Projected debt-to-income alone does not explain the harmful disbursement margin identified here. If the loan mainly harmed borrowers who were already visibly over-stretched on the lender's own affordability measure, the effect should rise with projected debt burden. It does not. The controls already account for credit quality, projected leverage, spending pressure, loan terms and purposes, family structure, applicant history, and employment and housing stability. Measured debt burden and the main applicant-level controls leave the harmful disbursement margin hidden before funds are sent.

B. Pre-Disbursement Observables Do Not Separate the Harmed Borrowers

Table 8 compares borrowers who complete loans on high-completion days with borrowers who complete them on low-completion days. Credit score, age, income, and projected debt-to-income have small standardized differences in means. The sharp difference is what happens later: subsequent default.

The comparison uses the standard complier-characterization logic. Early-week takers contain always-takers plus applicants induced to complete by faster processing. Late-week takers more closely approximate always-takers because the marginal applicants who are moved by higher friction do not complete. Under monotonicity, differences between these groups characterize the observable profile of the moved borrowers on average.

Table 8: Complier Characterization: Early- vs. Late-Week Takers

Variable	Mon–Tue takers	Fri–Sun takers	Difference	SMD
Credit score (0–100)	-0.172	-0.171	-0.001	-0.001
Age	0.062	0.059	+0.003	+0.003
Default rate (18mo)	12.1%	10.9%	+1.1pp	+0.035**
Loan amount (EUR)	12,737	12,589	+148	+0.010
Gross income (EUR)	3,475	3,563	-88	-0.002
Consolidation amount (EUR)	7,437	7,318	+119	+0.002
Proj. DTI (std.)	-0.007	-0.009	+0.003	+0.003
Expenses / ref. budget	0.055	0.083	-0.028	-0.028
Months at address	0.037	0.022	+0.015	+0.015

Notes: Early-week takers (Mon–Tue) include both always-takers and compliers; late-week takers (Fri–Sun) approximate always-takers, as compliers do not complete under higher operational friction. The difference in observables characterizes compliers (Angrist and Pischke 2009, Section 4.4.4). SMD = standardized mean difference (difference / pooled SD). Significance from Welch's t-test: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Estimated complier share: 1.2 percentage points (completion gap: Mon–Tue = 22.2%, Fri–Sun = 21.0%).

This is the paper's hardest screening fact. The main borrower and application variables available before disbursement do not separate harmed marginal borrowers from other marginal borrowers. Credit score, age, income, and projected debt-to-income move little across the completion margin; later default moves sharply after disbursement. Combined with the null debt-burden interaction, that fact explains why ex post non-default is such a misleading way to judge rejection quality.

The online appendix repeats the logic in the repeat-applicant sample used for the within-person design. Comparing the same applicant across weekdays leaves the main pre-disbursement observables moving little while later default moves. The main specification carries the evidentiary weight; the within-person estimate is a same-sign check.

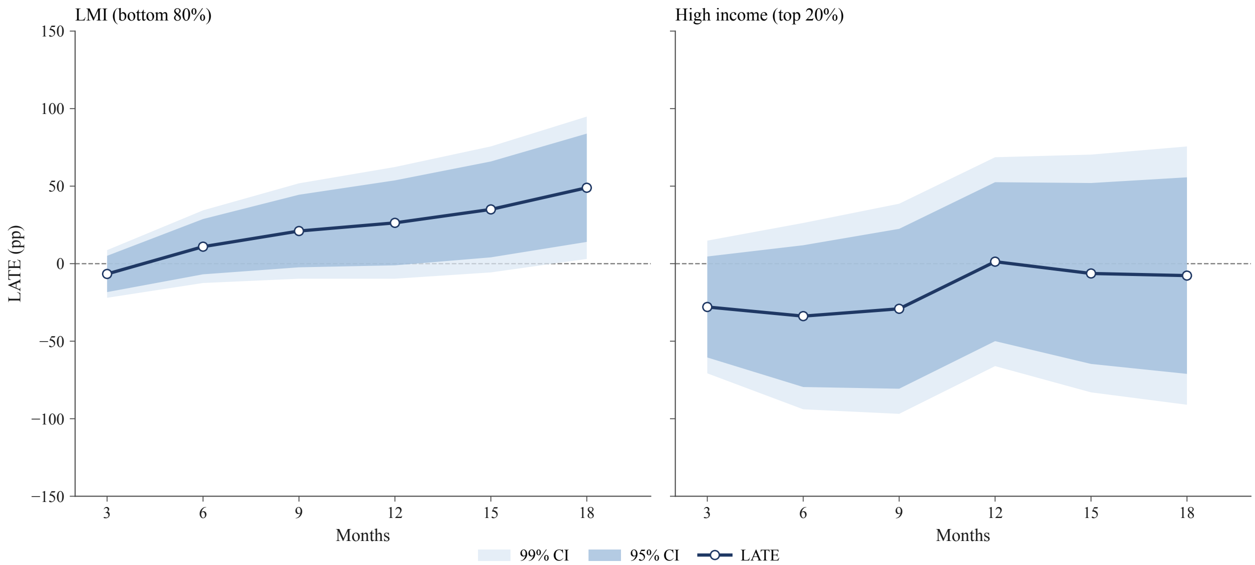

Figure 5: Treatment Effect by Income Group and Horizon

Notes: Separate instrumental-variables estimates by postal-code income group. Shaded regions show 95% confidence intervals. Standard errors are clustered at the borrower level.

C. The Harmful Margin Concentrates Where Delay Sensitivity Is Strongest

Columns 3 and 4 of Table 7 split the IV estimate by local-area income. This split is not another version of the continuous DTI test. Projected DTI is an applicant-level measured-affordability variable; the low- and moderate-income split is a postal-code proxy exercise for local resources and constraints that are not fully measured in the application file. At 18 months, the LMI split estimates a 49.0 percentage-point effect ($p = 0.006$; KP $F = 27.1$). The high-income split has a weak first stage (KP $F = 4.5$) and a -7.7 percentage-point estimate with $p = 0.812$, so it cannot test a comparable effect precisely. Figure 5 traces the split estimates over time.

The LMI split is where the paper detects the harmful disbursement margin and where weekday delays most strongly reduce completion. The high-income split cannot carry the same claim because weekday is a weak instrument there and the effect estimate is statistically indistinguishable from zero. Local-area income proxies for fallback capacity and local resource constraints that are not fully measured in the application file: weaker cash buffers, fewer movable assets, thinner support networks, unstable expenses, and behavioral bandwidth under repayment pressure not fully captured by credit score or projected DTI.

The local-area result is an incidence result, not a group-causality result. It identifies where repayment constraints are concentrated and where non-credit interventions are directly relevant: financial education, debt counseling, savings support, income-smoothing tools, or clearer risk disclosure at the point of offer. It does not imply that applicants in these areas are viable risks whose rejection is mistaken. That question is answered by the disbursement counterfactual. On the margin identified here, extending credit raises default.

D. Interpreting the Residual

The evidence establishes three facts. First, the effect does not rise with projected debt-to-income. Second, the IV identifies a harmful margin that the main application and platform variables do not separate before disbursement. Third, the area split detects that margin in LMI areas (49.0 percentage points; $p = 0.006$; KP $F = 27.1$), while the high-income split is weak-instrument and statistically insignificant (-7.7 percentage points; $p = 0.812$; KP $F = 4.5$).

These tests benchmark residual risk against commercial screening, not a thin researcher-chosen variable list. Approved applicants are borrowers whose underwriting profiles lenders judged acceptable. Within that pool, projected DTI does not move with default, the harm does not vary with projected DTI, and compliers and non-compliers are nearly identical on credit score, age, income, and projected DTI. The residual risk is residual to the screening commercial lenders have already applied.

The residual risk has two concrete components. One is fallback capacity: projected debt-to-income misses cash buffers, movable assets, support networks, and unstable expenses that make the same payment safer for one borrower than another. The other is behavioral bandwidth under repayment pressure: borrowers differ in how much room they have to absorb repayment strain and avoid compounding distress after funds arrive. The LMI split identifies where the risky margin is detected; the DTI and balance tests show that the main borrower-level variables summarized in the paper do not by themselves separate harmed marginal borrowers before funds are sent. The online appendix shows the same point descriptively: within broad credit-quality groups and projected-DTI bins, the LMI/high-income distinction is not mechanically reproduced by the standard borrower-level variables.

The timing evidence is direct. Processing friction moves completion in the direction the design requires. Applications submitted later in the week wait longer before cash reaches the borrower; relative to Monday, Friday applications take $+22.35$ more hours to reach payout. Weekday changes completion speed for a transaction already in motion, not the lender's offered terms.

The default timing profiles of weekday groups stay close once each group is normalized by its 18-month default rate. The online appendix reports interim default rates at 3, 6, 9, 12, and 18 months for Monday–Tuesday, Wednesday–Thursday, and Friday–Sunday applications. The maximum normalized spread is 2.36 percentage points, consistent with weekday mainly changing which applicants complete rather than the post-borrowing timing of recorded default within completed loans.

The result is a harmful disbursement margin detected in LMI areas, not separated by the main borrower-level variables, and activated by easier completion.

VII. Policy Implications

The policy point follows directly from the causal estimate. Once disbursement changes repayment outcomes, rejection quality cannot be judged by ex post repayment under rejection alone. The correct question is whether the applicant would still have repaid if the loan had been disbursed. On the margin identified here, the answer is no for many borrowers: disbursement raises 18-month default by 55.6 pp.

A. Why Ex Post Non-Default Is the Wrong Benchmark

Ex post non-default among rejected applicants compares the wrong states. On the weekday-shifted margin, borrowers who avoid default after rejection default when disbursement changes their debt path. The IV estimate shows this. Calling these rejections false rejections therefore misclassifies a causal effect as a screening mistake. The economic content is the sign of the intervention: the false-rejection interpretation pushes toward more disbursement on this margin, while the causal estimate says that extra disbursement creates default. For this group, non-default after rejection is not evidence that the lender should have lent.

The audit benchmark implicitly treats non-default without the loan as evidence of non-default with the loan:

$$E[Y_i(1) - Y_i(0) \mid \text{audit margin}] \approx 0.$$

The IV estimate rejects that assumption on the weekday-moved offered-applicant margin:

$$E[Y_i(1) - Y_i(0) \mid C] = 0.556.$$

Thus an applicant looks safe because the loan was not disbursed.

B. Incidence Is Not Causal Harm

Group-level approval gaps and default-prediction gaps are useful monitoring statistics, but they are not causal evidence that credit should have been extended. They show where repayment constraints are concentrated. They do not show whether the marginal applicant would have repaid under the loan. Credit allocation is made applicant by applicant: absent a binding portfolio constraint, a lender lends to every borrower whose expected repayment makes the loan viable. A lower approval rate for a coarse group therefore reflects the distribution of individual repayment constraints within that group, not evidence that profitable credit was denied.

The constructive use of incidence statistics is to locate where financial fragility is concentrated. They are descriptive, not causal. They guide credit education, debt counseling, savings support, and better disclosure. They should not substitute for the lending counterfactual. If policy treats non-default after rejection or group-level approval gaps as evidence that rejected applicants are viable risks, it pushes lending onto borrowers for whom disbursement itself causes default. If policy removes repayment-relevant local information, it first worsens prediction. That creates errors in both directions: bad marginal loans are disbursed, and viable marginal applicants are rejected. With capped upside and large default losses, both errors shrink viable lending. The contraction is not incidence-neutral: market size is set at the approval margin, so the first applicants rationed out are those closest to the break-even condition, including the lower-income or protected-status-correlated cells where screening constraints are most visible.

C. What the Result Changes

The paper identifies the object a planner needs first: the effect of disbursement on later distress for the borrowers moved by easier completion. For those offered applicants, the expected distress cost rises after the loan arrives. That matters for borrower welfare. It also matters for lenders, because, using observed lender-set contract terms and the opportunity-cost benchmark, one default absorbs the profit from roughly 6 fully repaid loans at the baseline 8.0% cost of capital and 5–10 loans over the 7.0–11.0% range. This is a debt-contract asymmetry, not an accounting footnote: lenders' contractual upside is capped, while default creates large downside losses.

That asymmetry changes the meaning of a false-rejection rule. Treating protective rejections as false rejections pushes credit to borrowers whose default risk rises because the loan is disbursed. The first effect is direct borrower harm. The appendix simulation translates that borrower-level result into portfolio stakes: under stated lender-response rules, forcing lending onto this margin lowers portfolio value and creates incentives to tighten elsewhere.

D. Direct Protection and Portfolio Stakes

The direct protective channel is the paper's core policy result. On the identified disbursement margin among offered applicants whose completion responds to weekday processing frictions, loan disbursement raises 18-month default by 55.6 percentage points. Rejection prevents those defaults. On this margin, disbursement imposes losses on both borrower and lender.

The repeat-applicant comparison is descriptive but stark: borrowers who reapply after an initial rejection and eventually proceed to disbursement default at 7.1% versus 14.0% for first-time approved applicants. That comparison supports the same screening interpretation without carrying the causal claim.

The portfolio logic follows from the payoff asymmetry. Once policy pushes lending to borrowers whose default risk rises after disbursement, the same capital produces more defaults and lower expected lender value. Under the stated lender-response exercise reported in the appendix, lenders restore portfolio value by tightening elsewhere; the exercise quantifies the incentive logic under maintained rules.

E. Why the Policy Margin Matters

The identified margin is narrow by design, and that is the point. The design captures the borrowers whose completion decision changes when the last step of borrowing becomes slightly easier or harder.

Digital-credit interventions operate exactly there: faster verification, fewer clicks, more insistent follow-up, softer cooling-off frictions, or performance targets that emphasize conversion.

The magnitudes are large in the units that matter. At the point estimate, one hundred extra completions induced on this margin generate about 55.6 additional 18-month defaults. Each default absorbs the profit from roughly 6 fully repaid average loans at the baseline cost of capital, with a 5–10 loan range under higher and lower funding-cost assumptions. The margin is local, but the losses are not small: a policy or platform change that raises completion for these borrowers destroys the surplus created by many otherwise performing loans. In lower-income areas, where the estimated effect is largest, this margin is especially consequential.

F. What Follows for Policy

The policy implication is to improve information and decision quality around the correct counterfactual. A centralized default registry, repayment-relevant local information, better real-time affordability information, clearer risk disclosure at the point of offer, cooling-off mechanisms that preserve deliberation, and market-based credit-building products address the screening problem directly. Preserving repayment-relevant statistical differentiation expands viable lending when it separates borrowers who would repay from borrowers whom the loan itself pushes into default. Better full-funnel training data have the same market-size logic: they help lenders identify viable borrowers inside high-incidence cells, which preserves surplus on safe loans and creates room to fund applicants closer to the margin. Mandates or conversion targets aimed at applicants who look safe only because the loan was never disbursed target the wrong margin.

The result is not an argument against credit access generally. It is an argument against treating ex post non-default without the loan as sufficient evidence that lending would have been safe. The result applies to the operational margin moved by faster verification, fewer completion steps, reminders, weaker cooling-off frictions, and conversion pressure.

The broader lesson is about modern digital credit markets and debt markets with capped upside and downside default risk. Smoother completion worsens welfare when the borrowers pulled in by the smoother process are the borrowers whose default risk rises after disbursement. Information restrictions work through the same margin: they make prediction worse near the approval cutoff, creating both bad disbursements and missed viable loans. The policy question asks whether the borrowers whose decisions are changed by easier completion or worse information are helped or harmed by loan disbursement. The evidence here answers that question on the identified margin: they are harmed, so treating the resulting protective rejections as false rejections targets the margin where disbursement creates default.

VIII. Conclusion

Ordinary lender records cannot answer whether a rejected applicant who avoids default after rejection would have defaulted after disbursement. Merging Finnish platform applications with credit-registry defaults for the full applicant pool, and using weekday processing delays to shift completion, lets the paper answer that counterfactual. On the identified disbursement margin among offered applicants whose completion responds to weekday processing frictions, loan disbursement raises 18-month default by 55.6pp. A proxy estimate puts their no-loan default rate at 11.3% and their implied with-loan rate near 66.9%. At the point estimate, 100 extra loan disbursements on this margin mean about 55.6 extra defaults. On this margin, rejection prevents defaults caused by disbursement.

The result changes the meaning of false rejection. The absence of default after rejection is not enough to show that a rejection was mistaken, because disbursement changes the outcome being used to judge the decision. Group incidence statistics are descriptive, not causal. They target support where repayment constraints concentrate, but they do not prove marginal rejected applicants are viable credit risks. For the applicants identified here, disbursement creates defaults that rejection prevents.

The payoff implication follows from the same fact. Debt contracts cap upside and expose lenders to large downside losses when borrowers default. Using observed lender-set contract terms and the opportunity-cost benchmark, one default absorbs the profit from roughly 6 fully repaid loans at the baseline 8.0% cost of capital, with a 5–10 loan range over plausible fintech funding costs. Treating these protective rejections as false rejections would direct lending toward borrowers whose default risk rises because of disbursement.

The paper also shows where pre-disbursement information carries the risk signal. Local-area income proxies for fallback capacity and local resource constraints that are not fully measured in the application file: the LMI split estimates a 49.0 percentage-point effect ($p = 0.006$; KP $F = 27.1$), while the high-income split has a weak first stage and a statistically insignificant -7.7 percentage-point estimate ($p = 0.812$; KP $F = 4.5$). Within the already-approved pool, credit score, age, income, and projected debt-to-income do not separate the borrowers whose completion is shifted by delays, even though later default differs sharply. The remaining risk is thinner buffers, weaker fallback options, and other traits not captured by credit score or projected debt-to-income alone after lender screening.

The contribution is to identify one operational margin on which easier loan disbursement is harmful and to show why that fact changes the policy object. The result is local, but it is local to the borrowers moved by faster verification, fewer completion steps, conversion pressure, and related interventions. The setting travels because the same debt-contract asymmetry holds outside this platform: competitive pricing that equalizes expected returns on average still leaves large realized downside risk and makes screening valuable. Once loan disbursement itself changes outcomes, policy should ask not only who looks safe without the loan, but also who is made worse off by disbursement.

Better policy should therefore improve information and decision quality around the correct benchmark. A centralized default registry, repayment-relevant local information, better real-time affordability information, clearer risk disclosure at the point of offer, cooling-off mechanisms that preserve deliberation, and market-based credit-building products are more consistent with the evidence than mandates based on ex post non-default. The constructive use of group-level evidence is to target support where financial fragility is concentrated, not to substitute incidence for causal evidence. Preserving repayment-relevant information improves prediction at the approval margin: fewer bad marginal loans are disbursed, and fewer viable marginal applicants are rejected. Treating protective rejections as false rejections harms portfolios through payoff asymmetry. The core result is simple: on the operational margin identified here, easier disbursement creates defaults that rejection prevents.

References

Adams, William, Liran Einav, and Jonathan Levin. 2009. "Liquidity Constraints and Imperfect Information in Subprime Lending." 252 citations (Crossref/DOI) [2025-10-23], *American Economic Review* 99, no. 1 (February 1, 2009): 49–84. ISSN: 0002-8282, accessed October 23, 2025. https://doi.org/10.1257/aer.99.1.49. https://pubs.aeaweb.org/doi/10.1257/aer.99.1.49.

Allcott, Hunt, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman. 2022. "Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending." *The Review of Economic Studies* 89, no. 3 (May 7, 2022): 1041–1084. ISSN: 0034-6527, 1467-937X, accessed August 18, 2023. https://doi.org/10.1093/restud/rdab066. https://academic.oup.com/restud/article/89/3/1041/6374508.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2022. "Consumer-lending discrimination in the FinTech Era." *Journal of Financial Economics* 143, no. 1 (January): 30–56. ISSN: 0304405X, accessed January 11, 2025. https://doi.org/10.1016/j.jfineco.2021.05.047. https://linkinghub.elsevier.com/retrieve/pii/S0304405X21002403.

Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri. 2020. "On the Rise of FinTechs: Credit Scoring Using Digital Footprints." Edited by Andrew Karolyi. 604 citations (Crossref/DOI) [2025-11-08], *The Review of Financial Studies* 33, no. 7 (July 1, 2020): 2845–2897. ISSN: 0893-9454, 1465-7368, accessed November 8, 2025. https://doi.org/10.1093/rfs/hhz099. https://academic.oup.com/rfs/article/33/7/2845/5568311.

DeFusco, Anthony A., Huan Tang, and Constantine Yannelis. 2022. "Measuring the welfare cost of asymmetric information in consumer credit markets." *Journal of Financial Economics* 146, no. 3 (December): 821–840. ISSN: 0304405X, accessed August 18, 2023. https://doi.org/10.1016/j.jfineco.2022.09.001. https://linkinghub.elsevier.com/retrieve/pii/S0304405X22001866.

Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. 2021. "Measuring Bias in Consumer Lending." *The Review of Economic Studies* 88, no. 6 (November 1, 2021): 2799–2832. ISSN: 0034-6527, accessed August 31, 2025. https://doi.org/10.1093/restud/rdaa078. https://doi.org/10.1093/restud/rdaa078.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2022. "Predictably Unequal? The Effects of Machine Learning on Credit Markets." *The Journal of Finance* 77, no. 1 (February): 5–47. ISSN: 0022-1082, 1540-6261, accessed September 18, 2024. https://doi.org/10.1111/jofi.13090. https://onlinelibrary.wiley.com/doi/10.1111/jofi.13090.

Hardt, Moritz, Eric Price, and Nathan Srebro. 2016. *Equality of Opportunity in Supervised Learning,* arXiv:1610.02413, October 7, 2016. Accessed August 29, 2025. https://doi.org/10.48550/arXiv.1610.02413. arXiv: 1610.02413[cs]. http://arxiv.org/abs/1610.02413.

Hau, Harald, Yi Huang, Chen Lin, Hongzhe Shan, Zixia Sheng, and Lai Wei. 2024. "FinTech Credit and Entrepreneurial Growth." *The Journal of Finance* 79, no. 5 (October): 3309–3359. ISSN: 0022-1082, 1540-6261, accessed January 10, 2025. https://doi.org/10.1111/jofi.13384. https://onlinelibrary.wiley.com/doi/10.1111/jofi.13384.

Karlan, Dean, and Jonathan Zinman. 2009. "Observing Unobservables: Identifying Information Asymmetries With a Consumer Credit Field Experiment." 290 citations (Crossref/DOI) [2025-10-23], *Econometrica* 77 (6): 1993–2008. ISSN: 0012-9682, accessed October 23, 2025. https://doi.org/10.3982/ECTA5781. http://doi.wiley.com/10.3982/ECTA5781.

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. 2018. "Algorithmic Fairness." *AEA Papers and Proceedings* 108 (May 1, 2018): 22–27. ISSN: 2574-0768, 2574-0776, accessed February 23, 2025. https://doi.org/10.1257/pandp.20181018. https://pubs.aeaweb.org/doi/10.1257/pandp.20181018.

Livshits, Igor, James C. Mac Gee, and Michèle Tertilt. 2016. "The Democratization of Credit and the Rise in Consumer Bankruptcies." *The Review of Economic Studies* 83, no. 4 (October): 1673–1710. ISSN: 0034-6527, 1467-937X, accessed July 25, 2024. https://doi.org/10.1093/restud/rdw011. https://academic.oup.com/restud/article-lookup/doi/10.1093/restud/rdw011.

Melzer, Brian T. 2011. "The Real Costs of Credit Access: Evidence from the Payday Lending Market*." *The Quarterly Journal of Economics* 126, no. 1 (February): 517–555. ISSN: 0033-5533, 1531-4650, accessed July 25, 2022. https://doi.org/10.1093/qje/qjq009. https://academic.oup.com/qje/article-lookup/doi/10.1093/qje/qjq009.

Meursault, Vitaly, Daniel Moulton, Larry Santucci, and Nathan Schor. 2022. *One Threshold Doesn't Fit All: Tailoring Machine Learning Predictions of Consumer Default for Lower-Income Areas.* Working paper (Federal Reserve Bank of Philadelphia) 22-39. Series: Working paper (Federal Reserve Bank of Philadelphia). Federal Reserve Bank of Philadelphia, November. Accessed September 18, 2024. https://doi.org/10.21799/frbp.wp.2022.39. https://www.philadelphiafed.org/-/media/frbp/assets/working-papers/2022/wp22-39.pdf.

Stiglitz, Joseph E., and Andrew Weiss. 1981. "Credit Rationing in Markets with Imperfect Information." *The American Economic Review* 71, no. 3 (June): 393–410. ISSN: 0002-8282, accessed August 29, 2025. https://www.jstor.org/stable/1802787.

Tantri, Prasanna. 2021. "Fintech for the Poor: Financial Intermediation Without Discrimination*." *Review of Finance* 25, no. 2 (March 1, 2021): 561–593. ISSN: 1572-3097, accessed August 30, 2025. https://doi.org/10.1093/rof/rfaa039. https://doi.org/10.1093/rof/rfaa039.

Zinman, Jonathan. 2010. "Restricting consumer credit access: Household survey evidence on effects around the Oregon rate cap." *Journal of Banking & Finance* 34, no. 3 (March): 546–556. ISSN: 03784266, accessed August 18, 2023. https://doi.org/10.1016/j.jbankfin.2009.08.024. https://linkinghub.elsevier.com/retrieve/pii/S0378426609002283.