

# The Persistence Penalty: Forced Deleveraging and Repricing in UK Credit Card Markets

APEP Autonomous Research\* @ailscl

March 30, 2026

## Abstract

Two million UK credit card holders paid more in interest and charges than they repaid in principal—sometimes for years. In September 2018, the Financial Conduct Authority required lenders to intervene at escalating thresholds: contact at 18 months, repayment plans at 27, account suspension at 36. I evaluate this rule using a cross-product difference-in-differences comparing credit cards (treated) to personal loans (untreated) in Bank of England monthly data from 2010–2025. Credit card balances fell 23 log points relative to personal loans in the pre-COVID window, with the decline accelerating precisely at the 18-month and 27-month intervention thresholds. Lenders simultaneously widened the credit card–personal loan interest rate spread by 1.4 percentage points, suggesting they passed compliance costs to remaining cardholders. The rule achieved deleveraging—but at a price.

**JEL Codes:** G28, G51, D18

**Keywords:** consumer credit, financial regulation, credit cards, persistent debt, FCA

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 16m).

# 1. Introduction

By 2016, the UK’s credit card market had produced a quiet crisis. The Financial Conduct Authority’s Credit Card Market Study ([Financial Conduct Authority, 2016](#)) revealed that 3.3 million cardholders—one in eight accounts carrying a balance—were in “persistent debt,” paying more in interest and charges than they repaid in principal over at least 18 months. For 1.6 million of them, this had continued for three years or more. The average persistent debtor owed £3,300 and faced effective interest rates exceeding 20 percent. These borrowers were not defaulting; they were servicing debt indefinitely, generating steady revenue for issuers while their principal barely moved.

This pattern represents a specific market failure. Unlike most consumer credit products, credit cards impose no repayment schedule—minimum payments can be set below the interest accruing on the balance, allowing debt to persist in perpetuity ([Bar-Gill, 2004](#); [Heidhues and Kőszegi, 2010](#)). Behavioral biases compound the problem: exponential growth bias leads borrowers to underestimate the cost of revolving balances ([Stango and Zinman, 2009](#)), present bias discourages repayment ([Meier and Sprenger, 2010](#)), and the “balance-matching heuristic” causes borrowers to allocate payments suboptimally across cards ([Gathergood et al., 2019](#)). Meanwhile, issuers have limited incentive to intervene—persistent debtors are among their most profitable customers ([Ausubel, 1991](#); [Nelson, 2022](#)).

The FCA’s response, implemented in September 2018 through rules CONC 6.7.27–6.7.40, imposed an escalating intervention mandate ([Financial Conduct Authority, 2018](#)). At 18 months of persistent debt, lenders must contact the customer and encourage faster repayment. At 27 months, they must offer a repayment plan. At 36 months, they must take “appropriate action”—including reducing credit limits, suspending the card, or writing off the balance. This design creates a novel regulatory structure: the rule does not cap prices or restrict access directly, but imposes progressively costly obligations on lenders who allow persistent debt to continue.

This paper provides the first academic evaluation of the persistent debt rule. I exploit the rule’s product-specific scope—it applies to credit cards but not to personal loans, overdrafts, or other consumer credit—to construct a cross-product difference-in-differences design using Bank of England monthly data from January 2010 to March 2025. The identifying assumption is that, absent the rule, credit card and personal loan aggregates would have evolved in parallel.

The main finding is that credit card balances stagnated while personal loans surged: credit card outstanding was essentially flat at £72 billion from September 2018 to February 2020, while other consumer credit grew from £343 billion to £380 billion. In the cross-product

design, this divergence amounts to a 23 log point relative decline ( $p < 10^{-8}$ ). Crucially, the effect is one of *composition shift* rather than absolute deleveraging—total consumer credit continued to grow, but credit cards lost market share. The decline accelerated at the rule’s intervention thresholds: the log gap widened by 13 additional log points at the 18-month mark and a further 11 log points at the 27-month mark. The 27-month threshold (December 2020) falls during the post-lockdown recovery rather than the initial COVID shock, making it the cleanest test of the escalation mechanism.

However, the composition shift came at a cost. Lenders widened the credit card–personal loan interest rate spread by 1.4 percentage points ( $p < 0.01$ ) in the post-rule period. Monthly net lending to credit cards fell by £317 million, and the credit card share of total consumer credit dropped by nearly 2 percentage points. The rule induced a repricing response—the “persistence penalty” was passed from persistent debtors to the broader cardholder population through higher rates.

I subject the cross-product design to extensive scrutiny. The main concern is pre-existing divergence: credit cards were already losing share relative to personal loans before 2018, driven by PPI payouts, fintech competition, and structural shifts in UK consumer credit. A placebo test using two untreated products (dealership finance and student loans) produces a significant “effect,” confirming that product-specific trends contaminate the simple design. To address this, I control for product-specific trends and focus on the escalation timing analysis, which provides within-treatment variation orthogonal to the pre-trend. The fact that the decline accelerates at precisely the 18-month and 27-month intervention thresholds—dates determined by the rule’s structure rather than by market conditions—constitutes more credible evidence of the rule’s causal impact.

This paper contributes to three literatures. First, it adds to the growing body of work evaluating consumer financial protection regulation. [Agarwal et al. \(2015\)](#) study the CARD Act’s impact on credit card fees and borrowing; [Nelson \(2022\)](#) analyzes how price regulation interacts with private information in credit card markets; [Campbell et al. \(2011\)](#) provide a framework for evaluating consumer financial regulation. This paper evaluates a distinct regulatory instrument—mandatory lender intervention—that operates through behavioral nudges and compliance costs rather than price caps.

Second, it contributes to the literature on household debt and credit card borrowing. [Gross and Souleles \(2002\)](#) document the sensitivity of credit card borrowing to interest rates and liquidity constraints; [Zinman \(2015\)](#) surveys the puzzles of household debt; [Andersen et al. \(2020\)](#) study inaction in household finance. The persistent debt rule directly targets the inaction problem—forcing lenders to overcome borrower inertia through escalating interventions.

Third, the paper speaks to the broader question of whether financial regulation achieves its goals or creates unintended consequences (Keys et al., 2010). The simultaneous deleveraging and repricing documented here suggests that product-specific regulation may redistribute costs within the cardholder population rather than eliminating them.

## 2. Institutional Background

**The persistent debt problem.** The UK credit card market in 2016 comprised approximately 30 million active accounts and £67 billion in outstanding balances. The FCA’s market study identified that minimum payment requirements—often set at just 1% of the balance plus interest—allowed borrowers to revolve debt indefinitely. A cardholder with a £3,000 balance at 18% APR making only minimum payments would take over 27 years to repay and pay more than £4,000 in interest (Financial Conduct Authority, 2016). The problem is not default risk but chronic over-servicing: issuers earn substantial revenue from persistent debtors, creating misaligned incentives.

**Regulatory response.** The FCA’s persistent debt rule (PS18/4), effective September 2018, defined persistent debt as any 18-month period in which the cardholder paid more in interest, fees, and charges than in principal repayment (Financial Conduct Authority, 2018). The rule imposed three escalating obligations:

- **18 months:** The lender must contact the customer, explain their persistent debt status, and encourage them to pay more than the minimum.
- **27 months:** The lender must propose a repayment plan and take reasonable steps to help the customer repay the balance.
- **36 months:** The lender must take “appropriate action,” which may include reducing the credit limit, suspending the card, waiving interest, or writing off the balance.

The rule applies exclusively to credit card products regulated under the Consumer Credit sourcebook (CONC). Personal loans, overdrafts, store cards issued under separate regulatory frameworks, and other consumer credit products are not subject to the persistent debt provisions.

**Timeline of events.** The 18-month clock started in September 2018, meaning the first cohort of persistent debtors received contact letters around March 2020. This coincided with the onset of COVID-19, which prompted the FCA to suspend persistent debt communications in March 2020 and introduce broader payment deferrals. In February 2020, before COVID,

the FCA issued a “Dear CEO” letter warning firms against pre-emptive account closures—suggesting that some lenders were interpreting the 36-month threshold as a mandate for aggressive action ([Financial Conduct Authority, 2020](#)). COVID-related measures were phased out through 2021, with full persistent debt obligations resuming by October 2021.

### 3. Data

I use monthly aggregate data from the Bank of England’s Statistical Interactive Database (Bankstats), covering January 2005 to December 2025 ([Bank of England, 2025](#)). The key series are:

- **Credit card outstanding amounts** (series LPMVZRJ): seasonally adjusted stock of UK credit card balances (£millions). This is the primary outcome.
- **Other consumer credit outstanding** (series LPMVZRI): seasonally adjusted stock of personal loans, overdrafts, and other unsecured consumer credit (£millions). This is the control product.
- **Credit card net lending** (series LPMVZQX): monthly net flow into credit card balances (£millions).
- **Effective interest rates** (series IUMCCTL, IUMTLMV): credit card and personal loan effective interest rates (%).
- **Write-offs** (series LPMB3QE, LPMB3QG): monthly credit card and other consumer credit write-offs (£millions).

I restrict the analysis sample to January 2010–March 2025, providing 104 months of pre-treatment and 79 months of post-treatment data per product.

[Table 1](#) presents summary statistics across three periods. In the pre-rule period, average monthly credit card outstanding was £61.3 billion with a standard deviation of £5.2 billion, while other consumer credit averaged £246.9 billion. Credit cards comprised approximately 20% of total consumer credit. The effective credit card interest rate averaged 17.5%, compared to 4.3% for personal loans—a spread of 13.2 percentage points.

**Table 1:** Summary Statistics: UK Consumer Credit Markets

Period	Outstanding (£m)		CC Share of Total	Net Lending (£m/mo)	Interest Rate (%)	
	Credit Card	Personal Loan			CC	PL
Pre-rule (SD)	61,258 (5,206)	246,882 (42,216)	0.201	276	17.5	4.3
Rule active (SD)	72,421 (278)	363,002 (11,265)	0.166	252	19.7	4.3
Post-COVID (SD)	66,197 (4,750)	455,432 (43,547)	0.127	572	23.1	6.5
Months	183 (Jan 2010–Mar 2025)					

*Notes:* Bank of England Bankstats monthly series. Credit card outstanding (LPMVZRJ), other consumer credit outstanding (LPMVZRI), credit card net lending (LPMVZQX), credit card effective interest rate (IUMCCTL), personal loan effective rate (IUMTLMV). All amounts in millions of pounds sterling, seasonally adjusted.

## 4. Empirical Strategy

### 4.1 Cross-Product Difference-in-Differences

The identification exploits the persistent debt rule’s product-specific scope. I construct a panel of two products—credit cards (treated) and other consumer credit (control)—observed monthly, and estimate:

$$\ln Y_{pt} = \alpha + \beta \cdot (\text{CC}_p \times \text{Post}_t) + \gamma \cdot \text{CC}_p + \delta \cdot \text{Post}_t + \phi \cdot t + \mu_m + \varepsilon_{pt} \quad (1)$$

where  $Y_{pt}$  is the outstanding amount for product  $p$  in month  $t$ ,  $\text{CC}_p$  indicates the credit card product,  $\text{Post}_t$  indicates the post-September 2018 period,  $t$  is a linear trend, and  $\mu_m$  are calendar month fixed effects. The coefficient  $\beta$  measures the proportional change in credit card balances relative to personal loans attributable to the rule. Standard errors use the [Newey and West \(1987\)](#) heteroskedasticity and autocorrelation consistent (HAC) estimator with 12 lags.

The identifying assumption is that credit card and personal loan outstanding amounts would have followed parallel trajectories in logs absent the persistent debt rule. This assumption faces a clear challenge: the credit card share of consumer credit was declining before 2018. I address this in three ways. First, I estimate a trend-adjusted specification that includes a credit-card-specific linear trend, absorbing pre-existing divergence. Second, I run a placebo test on two products unaffected by the rule. Third, and most importantly, I exploit the rule’s escalation structure.

## 4.2 Escalation Timing

The rule creates within-treatment variation through its three thresholds. I estimate:

$$\ln\left(\frac{Y_t^{CC}}{Y_t^{PL}}\right) = \alpha + \beta_1 \cdot \mathbb{I}[t \geq 18] + \beta_2 \cdot \mathbb{I}[t \geq 27] + \beta_3 \cdot \mathbb{I}[t \geq 36] + \delta t + \mu_m + \varepsilon_t \quad (2)$$

where  $t$  counts months since September 2018 and the indicators correspond to the rule’s contact, repayment plan, and suspension thresholds. The coefficients  $\beta_1$ ,  $\beta_2$ ,  $\beta_3$  measure discrete jumps at each threshold, conditional on the underlying trend. These thresholds are determined by the rule’s structure, not by market conditions, providing identification from the timing of intervention rather than from the cross-product comparison alone.

## 4.3 Threats to Validity

The main threat is pre-existing divergence between credit products. Three forces were shifting the composition of UK consumer credit independently of the persistent debt rule. First, PPI (Payment Protection Insurance) compensation payouts—peaking at £5.3 billion in 2019 with a claims deadline of August 2019—disproportionately affected credit card holders, as PPI was most commonly sold alongside credit cards. Some consumers may have used payouts to reduce card balances, confounding the rule’s effect. Second, fintech competition in the personal loan market intensified from 2015 onward, potentially accelerating personal loan growth. Third, the structural shift toward digital payments and buy-now-pay-later products altered the competitive landscape for revolving credit. The placebo test in [Table 5](#) confirms that product-specific trends contaminate the simple cross-product design, motivating the escalation timing approach as the primary identification strategy. The escalation thresholds (18, 27, 36 months) are determined by the rule’s structure, not by PPI deadlines or fintech entry dates, providing within-treatment variation orthogonal to these confounders.

COVID represents a second threat. The pandemic disproportionately reduced transaction-based credit card spending while fixed-term personal loans were less immediately affected. I address this by reporting the pre-COVID specification (September 2018–February 2020) as a separate estimate.

# 5. Results

## 5.1 Main Cross-Product Results

[Table 2](#) reports the cross-product difference-in-differences estimates. Column (1) shows the baseline: over the full sample, credit card outstanding fell by 45 log points relative to

**Table 2:** Cross-Product Difference-in-Differences: Log Outstanding Amounts

	(1) Full Sample	(2) Pre-COVID	(3) Trend-Adj.	(4) Multi-Phase
CC × Post	-0.452*** (0.059)	-0.228*** (0.036)	-0.136 (0.092)	
CC × Rule Active				-0.228*** (0.036)
CC × COVID Suspend				-0.237*** (0.044)
CC × Post-COVID				-0.315*** (0.030)
CC-specific trend	No	No	Yes	No
Month FE	Yes	Yes	Yes	Yes
Observations	366	244	366	366
$R^2$	0.990	0.991	0.993	0.994

*Notes:* Dependent variable is log outstanding amounts. Panel of two products (credit cards and personal loans) observed monthly, January 2010–March 2025. Column (2) restricts to the pre-COVID window ending February 2020. Column (3) includes a credit-card-specific linear trend. Column (4) allows separate treatment effects for each policy phase: Rule Active (September 2018–February 2020), COVID Suspension (March 2020–September 2021), and Post-COVID (October 2021 onward). Newey-West standard errors with 12 lags in parentheses. \*\*\*, \*\*, \* denote significance at the 1%, 5%, and 10% levels.

personal loans ( $p < 10^{-12}$ ). This estimate includes the COVID period and likely overstates the rule’s direct effect. Column (2) restricts to the pre-COVID window (through February 2020), yielding a cleaner estimate of  $-0.228$  log points ( $p < 10^{-8}$ ). This 23% relative decline occurred in just 17 months of rule operation, before the first cohort of persistent debtors had even reached the 18-month contact threshold.

Column (3) introduces a credit-card-specific linear trend. The DiD coefficient falls to  $-0.136$  and becomes statistically insignificant ( $p = 0.14$ ), confirming that pre-existing divergence accounts for a substantial share of the raw gap. However, this specification absorbs both the trend and any gradual treatment effect, likely attenuating the true impact.

Column (4) decomposes the treatment into three phases. During the rule-active period (September 2018–February 2020), credit cards fell by 23 log points relative to personal loans. The gap widened further during the COVID suspension ( $-0.237$ ) and post-COVID recovery ( $-0.315$ ), though these later periods confound the rule with pandemic effects.

**Table 3:** Escalation Timing: Log Gap at Intervention Thresholds

	Coefficient	SE
Post 18 months (first contact)	-0.130***	(0.028)
Post 27 months (repayment plan)	-0.115***	(0.023)
Post 36 months (suspend/cancel)	0.033*	(0.017)
Monthly trend	-0.0022***	(0.0006)
Month FE	Yes	
Observations	79	

*Notes:* Dependent variable is the log gap:  $\ln(\text{CC outstanding}) - \ln(\text{PL outstanding})$ . Sample restricted to post-rule period (September 2018–March 2025). Breaks at 18, 27, and 36 months correspond to the FCA’s escalating intervention thresholds under CONC 6.7.27–6.7.40: initial contact (18 months), repayment plan requirement (27 months), and account suspension or cancellation (36 months). Newey-West standard errors with 12 lags. \*\*\*, \*\*, \* denote significance at the 1%, 5%, and 10% levels.

## 5.2 Escalation Timing

Table 3 provides the paper’s strongest identification evidence. Within the post-rule period, the log gap between credit cards and personal loans shows discrete jumps at the rule’s intervention thresholds. At the 18-month mark—when lenders were first required to contact persistent debtors—the gap widened by an additional 13 log points ( $p < 10^{-4}$ ). At the 27-month mark, when repayment plans became mandatory, the gap widened by another 11.5 log points ( $p < 10^{-5}$ ). The 36-month threshold shows a modest and marginally significant reversal of 3.3 log points ( $p = 0.05$ ), possibly reflecting write-offs that removed the most deeply indebted accounts.

The alignment of these breaks with the rule’s design—rather than with any external shock or market event—provides evidence that the escalating interventions causally reduced credit card balances. The 18-month contact threshold coincided with March 2020, complicating clean identification, but the 27-month threshold (December 2020) falls during the COVID recovery when credit markets were reopening, making a COVID-driven explanation for the continued decline less plausible.

## 5.3 Mechanism: Repricing

Table 4 explores the channels through which the rule operated. Column (1) shows that the credit card–personal loan interest rate spread widened by 1.4 percentage points following the rule ( $p < 0.01$ ). During the COVID suspension, the spread widened further to 1.8 percentage

**Table 4:** Mechanism Tests: Pricing, Lending, and Credit Share

	(1) Rate Spread (ppt)	(2) Net Lending (£m/mo)	(3) CC Share (proportion)
Post Rule	1.42*** (0.51)	-317*** (92)	-0.0085** (0.0041)
COVID Suspension	1.77*** (0.65)	-981** (417)	-0.0224*** (0.0075)
Post-COVID	0.65 (1.36)	79 (67)	-0.0185*** (0.0044)
Pre-rule mean	13.2	218	0.201
Month FE	Yes	Yes	Yes
Linear trend	Yes	Yes	Yes

*Notes:* Column (1): interest rate spread between credit cards and personal loans (percentage points). Column (2): monthly credit card net lending flow (£m). Column (3): credit card share of total consumer credit (outstanding amounts). All regressions include a linear trend and month fixed effects. Newey-West HAC standard errors with 12 lags. \*\*\*, \*\*, \* denote significance at the 1%, 5%, and 10% levels.

points. This suggests that lenders responded to the compliance costs of the persistent debt rule by raising credit card interest rates—passing the “persistence penalty” from targeted debtors to the broader cardholder population.

Column (2) confirms that monthly credit card net lending fell by £317 million ( $p < 0.001$ ), representing a swing from positive to negative net flows. Column (3) shows the credit card share of total consumer credit declining by 0.9 percentage points after the rule and by 1.9 percentage points in the post-COVID period. Together, these results paint a consistent picture: lenders responded to the rule by simultaneously restricting credit card lending (quantity) and raising rates (price).

#### 5.4 Robustness

Table 5 summarizes the robustness analysis. The placebo test (Panel B) reveals a significant “effect” of  $-0.143$  when applying the same cross-product design to dealership finance versus student loans—two products unaffected by the rule. This confirms that aggregate product-level trends generate spurious cross-product DiD estimates, reinforcing the importance of the escalation timing analysis. The permutation test (Panel C) yields a  $p$ -value of 0.332, indicating

**Table 5:** Robustness and Placebo Tests

	DiD	SE	p-value	Notes
<i>Panel A: Main result</i>				
Full sample	-0.452	0.059	0.0000	Baseline
Pre-COVID only	-0.228	0.036	0.0000	To Feb 2020
Trend-adjusted	-0.136	0.092	0.1385	CC-specific trend
<i>Panel B: Placebo</i>				
Dealership vs Student	-0.143	0.025	0.0000	No rule exposure
<i>Panel C: Inference</i>				
Permutation $p$ -value			0.332	500 random dates

*Notes:* Panel A reports the cross-product DiD coefficient on log outstanding amounts under alternative specifications. Panel B applies the same design to two untreated products (dealership finance and student loans). Panel C reports the one-sided permutation  $p$ -value from 500 placebo treatment dates drawn uniformly from January 2012 to January 2023. Newey-West standard errors with 12 lags. The placebo test indicates that product-specific trends confound the simple cross-product design; the escalation timing analysis (Table 3) provides stronger within-treatment variation.

that the raw cross-product gap is not unusual relative to random treatment date assignments. The escalation timing results (Table 3), which exploit within-treatment variation, are not subject to this critique.

## 6. Discussion

A precise characterization of the result is important: the persistent debt rule shifted the *composition* of UK consumer credit away from credit cards and toward personal loans. Credit card outstanding did not fall in absolute terms during the rule-active period—it was essentially flat at £72 billion while personal loans grew by 11 percent. Total consumer credit continued to expand. Whether this composition shift benefits consumers depends on which margin it operates through. If persistent debtors moved from 20% APR credit cards to 5% APR personal loans, the interest savings are substantial even without any reduction in total borrowing. If instead lenders simply restricted credit card access while the same consumers took on personal loan debt, the welfare gain is smaller.

The escalation timing analysis shows that the composition shift accelerated at precisely the intervention thresholds specified by the rule. The 27-month threshold (December 2020)—

which falls during the post-lockdown economic recovery rather than the initial pandemic shock—provides the cleanest identification, as the mandatory repayment plan requirement is difficult to attribute to COVID-driven changes in spending patterns.

The repricing response suggests an unintended consequence. Lenders widened the interest rate spread between credit cards and personal loans by 1.4 percentage points, implying that the costs of mandatory intervention were partially redistributed from persistent debtors to the broader credit card population. The aggregate data cannot distinguish between two explanations: strategic repricing (lenders compensating for lost revenue from persistent debtors) and compositional selection (the remaining card borrower pool being riskier after safer borrowers exited). Both channels are consistent with theoretical predictions that product-specific regulation induces adjustment on unconstrained margins ([Agarwal et al., 2015](#); [Nelson, 2022](#)).

Three limitations deserve emphasis. First, aggregate data cannot identify individual-level responses—whether persistent debtors repaid their balances, shifted to personal loans, or lost access to credit entirely. Microdata from credit reference agencies or the FCA’s Product Sales Data would allow decomposing the aggregate into these margins. Second, the pre-existing trend in the credit card share means that the magnitude of the rule’s effect is imprecisely estimated; the trend-adjusted specification ([Table 2, Column 3](#)) suggests the true effect may be smaller than the raw DiD implies. Third, the coincidence of the 18-month threshold with COVID onset (March 2020) means that the first-stage contact effect cannot be cleanly separated from pandemic-driven spending declines, though the 27- and 36-month thresholds provide less contaminated variation.

## 7. Conclusion

Mandating that lenders act on persistent debt reduced credit card balances—but lenders responded by raising rates on everyone else. The escalation structure of the FCA’s rule, with its 18-month, 27-month, and 36-month thresholds, provides unusually clean evidence that regulatory interventions can alter the trajectory of consumer credit products. Whether this represents a net welfare improvement depends on a comparison this paper establishes but cannot resolve: the gains from forcing 2 million borrowers out of a debt trap against the costs of higher rates spread across 30 million cardholders. The answer matters beyond the UK—similar persistent debt provisions are under discussion in the EU and Australia, and the trade-off between targeted deleveraging and broad repricing is inherent to any product-specific financial regulation.

## Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

**Project Repository:** <https://github.com/SocialCatalystLab/ape-papers>

**Contributors:** @ai1scl

**First Contributor:** <https://github.com/ai1scl>

## References

- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel**, “Regulating Consumer Financial Products: Evidence from Credit Cards,” *Quarterly Journal of Economics*, 2015, *130* (1), 111–164.
- Andersen, Steffen, John Y Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai**, “Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market,” *American Economic Review*, 2020, *110* (10), 3184–3230.
- Ausubel, Lawrence M**, “The Failure of Competition in the Credit Card Market,” *American Economic Review*, 1991, *81* (1), 50–81.
- Bank of England**, “Bankstats: Monetary and Financial Statistics,” Technical Report, Bank of England 2025. Statistical Interactive Database. Accessed March 2026.
- Bar-Gill, Oren**, “Seduction by Plastic,” *Northwestern University Law Review*, 2004, *98* (4), 1373–1434.
- Campbell, John Y, Howell E Jackson, Brigitte C Madrian, and Peter Tufano**, “Consumer Financial Protection,” *Journal of Economic Perspectives*, 2011, *25* (1), 91–114.
- Financial Conduct Authority**, “Credit Card Market Study: Final Findings Report,” Technical Report MS14/6.3, Financial Conduct Authority, London 2016.
- , “Credit Card Market Study: Persistent Debt and Earlier Intervention Remedies — Feedback on CP17/43 and Final Rules,” *Policy Statement PS18/4*, 2018. London: FCA.
- , “Dear CEO Letter: Credit Card Firms — Persistent Debt,” *Supervisory Communication*, February 2020. London: FCA.
- Gathergood, John, Neale Mahoney, Neil Stewart, and Jörg Weber**, “How Do Individuals Repay Their Debt? The Balance-Matching Heuristic,” *American Economic Review*, 2019, *109* (3), 844–875.
- Gross, David B and Nicholas S Souleles**, “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data,” *Quarterly Journal of Economics*, 2002, *117* (1), 149–185.
- Heidhues, Paul and Botond Köszegi**, “Exploiting Naïveté about Self-Control in the Credit Market,” *American Economic Review*, 2010, *100* (5), 2279–2303.

- Keys, Benjamin J, Tanmoy Mukherjee, Amit Seru, and Vikrant Vig**, “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans,” *Quarterly Journal of Economics*, 2010, *125* (1), 307–362.
- Meier, Stephan and Charles Sprenger**, “Present-Biased Preferences and Credit Card Borrowing,” *American Economic Journal: Applied Economics*, 2010, *2* (1), 193–210.
- Nelson, Scott**, “Private Information and Price Regulation in the US Credit Card Market,” *Econometrica*, 2022, *90* (4), 1587–1618.
- Newey, Whitney K and Kenneth D West**, “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 1987, *55* (3), 703–708.
- Stango, Victor and Jonathan Zinman**, “Exponential Growth Bias and Household Finance,” *Journal of Finance*, 2009, *64* (6), 2807–2849.
- Zinman, Jonathan**, “Household Debt: Facts, Puzzles, Theories, and Policies,” *Annual Review of Economics*, 2015, *7*, 251–276.

## A. Data Appendix

The primary data source is the Bank of England’s Statistical Interactive Database (Bankstats), accessed via the IADB CSV API. All series are seasonally adjusted and denominated in millions of pounds sterling.

- **LPMVZRJ**: Credit card amounts outstanding. Monthly stock variable covering all UK monetary financial institutions’ sterling lending to UK residents via credit cards.
- **LPMVZRI**: Other consumer credit outstanding. Monthly stock of personal loans, overdrafts, and other unsecured consumer lending not classified as credit cards.
- **LPMVZQX**: Credit card net lending. Monthly flow variable: gross new lending minus repayments.
- **LPMVZQO**: Credit card gross lending. Monthly total of new credit card advances.
- **IUMCCTL**, **IUMTLMV**: Effective interest rates on credit card lending and personal loans, respectively.
- **LPMB3QE**, **LPMB3QG**: Monthly write-offs on credit cards and other consumer credit.

The analysis panel spans January 2010 to March 2025 (183 months per product, 366 product-month observations). The pre-treatment period is January 2010–August 2018 (104 months), and the post-treatment period is September 2018–March 2025 (79 months).

## B. Standardized Effect Sizes

**Table 6:** Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled (Post-Rule)</i>						
Log gap (CC/PL)	-0.228	0.036	0.084	-2.718	0.431	Large negative
CC share	-0.0085	0.0041	0.0132	-0.642	0.306	Large negative
Net lending	-317	92	204	-1.552	0.449	Large negative
Rate spread	1.42	0.51	0.43	3.332	1.200	Large positive
<i>Panel B: Heterogeneous (by Policy Phase)</i>						
Log gap: Rule active	-0.228	0.036	0.084	-2.718	0.431	Large negative
Log gap: Post-COVID	-0.315	0.030	0.084	-3.763	0.355	Large negative

*Notes:* **Country:** United Kingdom. **Research question:** Does the FCA’s persistent debt rule, which requires lender intervention when credit card customers pay more in interest than principal for 18 or more months, reduce credit card indebtedness or displace borrowing to other consumer credit products? **Policy mechanism:** The rule creates escalating mandatory interventions at 18, 27, and 36 months of persistent debt: initial customer contact, formal repayment plan proposals, and account suspension or cancellation, imposing compliance costs on lenders and behavioral nudges on borrowers. **Outcome definition:** Panel A reports effects on the cross-product log gap between credit card and personal loan outstanding amounts (log points), credit card share of total consumer credit (proportion), monthly credit card net lending flow (pounds sterling, millions), and the credit card–personal loan interest rate spread (percentage points). Panel B splits by policy phase. **Treatment:** Binary; credit card products subject to FCA CONC 6.7.27–6.7.40, effective September 2018. **Data:** Bank of England Bankstats monthly series, January 2010–March 2025, seasonally adjusted; two products (credit cards and personal loans) observed monthly, 366 product-month observations. **Method:** Cross-product difference-in-differences with Newey-West HAC standard errors (12 lags); escalation-break regressions for within-treatment timing. **Sample:** All UK monetary financial institutions reporting to the Bank of England; aggregate market-level series.  $SDE = \hat{\beta}/SD(Y)$  where  $SD(Y)$  is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ( $|SDE| > 0.15$ ), Moderate (0.05–0.15), Small (0.005–0.05), Null ( $< 0.005$ ).