NBER WORKING PAPER SERIES

INTEREST RATE CAPS, COMPETITION, AND STRATEGIC BORROWING: EVIDENCE FROM KENYA

Aroon Narayanan
Tavneet Suri
Prashant Bharadwaj

Working Paper 35166
http://www.nber.org/papers/w35166

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

We are grateful to seminar audiences at Berkeley Haas School of Business, Columbia University, MIT Sloan, The World Bank, Microsoft Research, MIT, Michigan University, George Washington University, and the University of Washington for comments. No financial support was needed for this project. We worked with the Commercial Bank of Africa to get data access from their M-Shwari platform under an NDA. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Interest Rate Caps, Competition, and Strategic Borrowing: Evidence from Kenya
Aroon Narayanan, Tavneet Suri, and Prashant Bharadwaj
NBER Working Paper No. 35166
May 2026
JEL No. G51, L13, O55

ABSTRACT

We study Kenya's 2016 interest-rate regulation, which capped bank lending rates but left one digital platform, called M-Shwari, exempt on the lending side while imposing a deposit-rate floor across all lenders in the market. Using borrower-level administrative data, survey data, and an RD around the implementation date, we show three main results. First, lending on the exempt platform rose, but with the safest borrowers substituting away toward cheaper capped credit. Second, riskier borrowers increase their savings to build up their credit limits. Third, on the supply side, M-Shwari raises the limits for the safest borrowers in an attempt to retain them. We build and estimate a simple model of screening and credit limit-setting to interpret these reallocations and compute welfare. The observed carve-out for M-Shwari preserves access for high-risk borrowers but yields a slight aggregate welfare decline relative to pre-policy. However, a uniform (across all lenders) interest rate cap counterfactual generates substantially larger welfare losses by entirely eliminating credit for high-risk borrowers.

Aroon Narayanan
Massachusetts Institute of Technology
aroon@mit.edu

Tavneet Suri
Massachusetts Institute of Technology
MIT Sloan School of Management
and NBER
tavneet@mit.edu

Prashant Bharadwaj
University of California, San Diego
Department of Economics
and NBER
prbharadwaj@ucsd.edu

1 Introduction

Credit markets are widely viewed as central to economic growth (Levine and Zervos, 1998), yet there is persistent debate on whether—and how—these markets should be regulated (Benston and Kaufman, 1996; Campbell et al., 2011). Consumer-protection concerns often motivate regulation, especially where market power or information frictions restrict access for otherwise creditworthy borrowers. At the same time, tighter rules can render some market segments unprofitable, leading credit providers to raise collateral or screening thresholds and exclude precisely those borrowers regulation aims to help. These tradeoffs are particularly salient with interest rate caps, a common policy tool used across many countries and currently being discussed in the United States as well (Maimbo and Henriquez Gallegos, 2014; Ferrari, Masetti, and Ren, 2018).

Despite their prevalence — nearly half of countries in Sub-Saharan Africa and Latin America maintain some form of cap on interest rates (Maimbo and Henriquez Gallegos, 2014) — measuring the behavioral and distributional consequences of these caps is challenging. In many developing economies, informal credit operates outside formal regulation and is therefore exempt from such regulation, and caps are frequently applied uniformly across lenders, limiting the scope to observe any substitution or competitive responses across lenders. Addressing these challenges requires a setting where (i) formal access is widespread and low-cost, and (ii) the policy differentially affects competing lenders.

We study Kenya's 2016 interest rate regulation, which provides precisely such a setting. The law had two components. On the lending side, it capped bank loan rates at four percentage points above the Central Bank of Kenya (CBK) base rate — which at enactment stood at 8.6%, pinning the lending cap at 14.5%. This was well below prevailing market rates: banks had been charging around 18% on secured loans, and considerably more on unsecured credit. On the deposit side, the law simultaneously imposed a floor requiring banks to pay savers at least 70% of that same base rate, or roughly 6.2%.

These regulations were partly applied asymmetrically, in that they were not applied to non-bank financial institutions so informal banking and credit were exempt as were organizations providing only credit. In addition, on the lending side, the consumer digital lending platform, called M-Shwari and operated by (at the time) the Commercial Bank of Africa (CBA) at the time was exempt because the fee it charged consumers was called a facilitation fee and not an interest rate when they got approval for the product from the CBK. However, competing digital platforms at other banks were not exempt. In addition, importantly, the deposit floor regulation applied to everyone, including M-Shwari, so savings returns did not differ across any of these platforms. M-Shwari was, at the time, the largest digital banking product in the country, though with competitors, all tightly integrated with M-PESA (Kenya's popular mobile money platform) and all offering remunerated savings and small, uncollateralized 30-day loans. By 2017, before the regulation, over 80% of mobile phone owners had an M-Shwari account, about 30% of them had borrowed on M-Shwari (Figure A1).

In this paper, we analyze the impact of the lending interest rate caps on borrower and bank behavior, and, using insights from a simple model and structural estimation, study their welfare consequences. First, using administrative (daily level) data and a regression discontinuity (RD) around the reform date, we provide micro-level reduced form evidence on how borrowers sort in response to the rate caps. Lending on the exempt (and more expensive) platform rises: the daily probability of having an M-Shwari loan increases by roughly 10% off a 1.8% baseline, and loan amounts rise by about 1.1%. New access expands as well: the probability of taking a first M-Shwari loan increases by about 70%. These shifts are not uniform across risk types. Borrowers with higher initial credit scores are significantly less likely to borrow on M-Shwari post-caps and reduce loan amounts, consistent with substitution toward cheaper capped products. At the same time, we see that M-Shwari raises credit limits—especially for safer customers (an additional score point is associated with a 15% increase in the limit)—suggesting a supply-side attempt to retain these safer customers. On the savings margin, where balances help build limits, overall saving on M-Shwari increases after the reform; the rise is concentrated among riskier customers, consistent with strategic saving to regain borrowing capacity.

Second, we formalize these patterns in a simple model of competition in digital credit platforms with heterogeneous repayment risk. Two lenders set screening thresholds (as a function of credit scores) and credit limits (one being the exempt platform, M-Shwari, and the second being the competitor). The rate caps reduce the competitor's price but tighten its screening cutoff, excluding marginal borrowers. Because M-Shwari's price is unchanged, excluded borrowers flow to them. In competitive segments (where both lenders remain active), M-Shwari has an incentive to raise limits to retain high credit score customers who can access cheaper capped loans elsewhere; in segments where their competitors effectively exit, M-Shwari may reduce limits as competitive pressure declines. Borrowers can invest in raising their firm-specific credit limits (e.g., by saving), and the common deposit floor lowers the cost of doing so across all platforms equally.

Third, we bring the model to the data. We combine the administrative panel data with survey data (from Suri, Bharadwaj, and Jack, 2021) and estimate the structural parameters of the model that govern screening and credit limit-setting. We need the model to assess welfare, because reduced-form shifts in volumes and sorting do not map cleanly into consumer surplus when competition weakens in some markets and strengthens in others. The estimated model rationalizes the RD facts and lets us compute borrower welfare and run counterfactuals. These exercises show that the observed price carve-out for M-Shwari preserved access for riskier borrowers (albeit at higher prices) and yielded a small aggregate welfare decline relative to the pre-policy baseline, while a uniform rate cap would have sharply reduced welfare by eliminating credit for low credit score borrowers altogether.

Related to our paper is work on interest rate ceilings in the U.S. payday and small-dollar credit markets (e.g., Temin and Voth, 2008; Zinman, 2010; Benmelech and Moskowitz, 2010; Melzer and Schroeder, 2017; among others). Across settings, caps tend to impact access and the provision of loans. In Oregon, the cap reduced payday loan use and is linked to worse short-run outcomes among affected households (Zinman, 2010). When Rhode Island tightened its cap, borrowing sequences lengthened and defaults rose (Fekrazad, 2020). On the Prosper online lending marketplace, higher state usury ceilings nudged marginal loans over the line without clear default effects (Rigbi, 2013). Illinois's 36% all-in cap shifted credit toward prime borrowers (Bolen, Elliehausen, and Miller, 2023). Finally, state interest-rate caps severely reduce the presence of payday lending establishments (Dasgupta and Mason, 2020). These studies underscore the distributional impacts of interest rate limits in developed-country settings with well-developed financial markets.

In developing economies with different institutional environments (more informal credit, less competition in the financial sector, more missing markets, etc.), studies show largely negative impacts of interest rate caps on riskier consumers and firms. As these tend to be precisely the households and firms that might benefit the most from credit (Grandolini, 2015), the overall welfare consequences of such caps in developing countries could be much greater. For Chile, Schmukler et al. (2022), Madeira (2019), and Cuesta and Sepulveda (2021) examine nationwide tightening of usury ceilings between 2013 and 2015 and document substantial rate compression alongside reductions in credit, with losses concentrated among riskier borrowers. Cuesta and Sepulveda (2021) further build a structural model to assess welfare, estimating average consumer-surplus declines on the order of 2.5% of annual income, with larger losses at the bottom of the risk distribution. Studying the same interest rate caps as we do in Kenya, Safavian and Zia (2018) and Alper et al. (2020) use bank-level panels and qualitative evidence to show sharp declines in aggregate lending—especially to SMEs—and retrenchment from higher-risk segments. Burga, Nivin, and Yamunaqué (2022) find similar evidence when Peru introduced interest rate caps in 2021. Recent work by Kuroishi, LaPoint, and Miyauchi (2025) offers a different perspective by studying corporate loan interest rate caps applied to large firms in Bangladesh. Notably, this cap did not apply to small and medium-sized firms. Under this particular policy, they find banks *increased* lending to firms and did not ration credit to riskier borrowers.

Our study complements the extant work in the developing world by leveraging borrower-level, day-by-day administrative data with original credit scores and credit limits, and by analyzing an asymmetric reform: Kenya's lending-side carve-out for M-Shwari, not for its direct competitors, alongside a uniform deposit-rate floor. This asymmetry is the key innovation here: we can leverage it to estimate both demand and supply sides of the digital credit market and analyze the competitive response on the part of exempt and non-exempt lenders. The RD evidence shows how competitors' tighter screening reallocated marginal (riskier) borrowers toward the exempt platform, while safer borrowers substituted to cheaper capped credit; importantly, we also document strategic saving to rebuild limits and active limit adjustments for less risky borrowers by M-Shwari. A simple model and structural estimation then map these reallocations into welfare: under the observed carve-out, aggregate welfare falls modestly while access for high-risk borrowers is preserved (albeit at worse terms), whereas a uniform-cap counterfactual produces substantially larger welfare losses concentrated among low-score borrowers. In this way, our results align with the Chilean evidence on the price–access trade-off while highlighting how policy design and digital-platform competition shape distributional impacts.

2 Background on the Digital Banking Sector and Interest Rate Caps

Kenya has seen an explosion of digital banking in recent years. Our analysis focuses on the first and largest digital banking platform, M-Shwari, and its main direct competitors (i.e. other digital lenders). In 2012, CBA launched M-Shwari, a fully digital bank account operating over the rails of mobile money, with remunerated savings and credit. M-Shwari is not a subsidiary of CBA but simply a digital product (and platform) launched by CBA. The M-Shwari account is linked to M-PESA, the popular mobile money service in Kenya. It is opened and operated through a USSD application on an account holder's phone. Loans are dispensed into users' mobile money accounts and the money can be transferred from there into and out of their M-Shwari accounts at no cost. The withdrawals and deposits out of M-Shwari accounts use the existing mobile money infrastructure in the country. M-Shwari offers both a remunerated savings account and the opportunity to get short term loans.

The take up of M-Shwari has been remarkable: within two years of its launch there were over 4.5 million active users (20% of adults) and approximately 10 million accounts had been opened. M-Shwari is credited with making CBA a major player in the lending market and has been an important source of competitive advantage for CBA. As of 2017, CBA had over 50% of the loan accounts and just under 50% of deposit accounts in the country (Central Bank of Kenya 2017; Suri and Gubbins 2018). At the time of the imposition of the interest rate caps, over 80% of mobile phone owners in Kenya had an account on M-Shwari (Figure A1 and

Figure A2).

A major draw for signing up for M-Shwari is the loan product where approved customers have access to small, short-term, 30 day, 7.5% fee loans, without any credit history. The loans disbursed by M-Shwari are uncollateralized and start at small amounts, with the first loan sometimes as low as KSh 100 ($1) but sometimes as high as KSh 10,000. Over time, even if an individual starts off with a low credit limit, if they repay and (importantly) save on M-Shwari, they can grow their limit. Behind the loan approval process is a set of credit approval and scoring rules based on data from the user's mobile money and phone records (see Suri, Bharadwaj, and Jack, 2021).

The credit scoring process gives individuals a loan limit which increases upon the timely repayment of a loan as well as through accumulating savings in M-Shwari.[1] A customer's initial credit score on M-Shwari is based on the customer's history of M-PESA transactions and mobile phone expenditures and this credit score is mapped onto a credit limit, including a zero limit for those not approved. Note that this credit score is built by M-Shwari and is not a common credit score that is available to all banks. It uses non-traditional data on mobile phones and mobile money from an agreement with Safaricom, the main cell phone company in Kenya that also built M-PESA. Over time, the customer's credit limit evolves based on their interactions with M-Shwari, by either borrowing and repaying (if they default, their credit limit reverts to zero) or by saving in their M-Shwari accounts.

Importantly for our paper, the bank never updates the initial credit scores, which are assigned automatically and immediately upon account opening, irrespective of whether a loan is requested by the customer. Customers are not told their credit scores, only their credit limit. However, even customers not initially approved for a loan can be approved for one in the future by saving on M-Shwari. Savings are used to update their credit *limits* even though their credit *scores* in the system remain unchanged. For those initially not approved for loans, saving on M-Shwari is therefore the key way to get a positive credit limit in the future. On the savings side, before the interest rate regulations, the annual interest rates on savings were between 2% and 5%, depending on the balance in the account.

A few years after M-Shwari launched, Kenya Commercial Bank (KCB) and Equity Bank launched similar products. For KCB, the data used to score customers is the same as that supplied for M-Shwari. There has also been an influx of purely digital lenders that do not take deposits but offer loans (and hence are not banks). The most popular such lenders, Tala and Branch, use a smart phone application, which scrapes the phone for data on mobile money transactions and phone records, hence use similar data too. Figure A1 illustrates the market,

[1] If a loan is not paid on time, the loan is extended for another 30 days with a 7.5% fee charged on the outstanding balance. After 120 days of non-payment, the borrower is reported to the credit reference bureau.

showing the fraction of individuals that had a digital loan across different companies in 2017 (see Totolo (2018)). M-Shwari had the majority market share, but KCB and Equity's products were close competitors and catching up. Figure A2 (see Suri and Gubbins (2018)) highlights how over 80% of mobile phone owners had an M-Shwari account in 2017.

2.1 Interest Rate Regulation

The Kenyan parliament imposed stringent caps on interest rates on September 14, 2016. The lending rate cap was set at four percentage points above the base rate published by the Central Bank of Kenya (CBK) and a deposit rate floor at 70% of the same base rate. At the time the law went into effect, the six largest banks in Kenya (which together held 87% of all deposit accounts and 76% of all loan accounts) charged interest rates on loans far higher than the cap. The average annual lending interest rate was 18% (the rate only applied to collateralized lending as there was no uncollateralized lending at this rate) and the 91-day Treasury bill rate was 8.6%. While the policy allowed for changes to the interest rate cap, at the time of implementation, the lending cap was set at 14.5%, far below prevailing market rates.

These regulations were proposed as part of the Banking (Amendment) Act 2016. Once the Act was passed in Parliament, the President signed it into law on 24 August 2016 and it went into effect on September 14. Though the law had been debated in parliament and discussed in the press before it went into effect, the President signing the bill was a shock to the industry and the economy. There was a strong sense that the President would veto the bill if it passed Parliament since the President's family held the largest share (approximately 25%) in CBA; hence, he would be hurting his own interests by signing the law. There is a short time window between 24 August and 14 September. We do not see strong evidence of anticipation effects and, if anything, we see a small reduction in loans, given it was known that interest rates would fall in a matter of days.

These legal restrictions on interest rates were imposed on banking institutions but non-bank financial institutions (such as savings and credit cooperatives, microfinance institutions, non-deposit taking lenders) were exempt. From Figure A1, companies like Tala and Branch were exempt since they are non-bank financial institutions (i.e. non deposit taking). Traditional banks like Equity and KCB were not exempt; the caps applied on their digital credit platform as well. Interest rates on KCB's digital credit product changed from 6% a month to 1.1% and on Equity Bank's digital product from as high as 9% to 1.1% overnight.

M-Shwari, however, was exempt from the interest cap on lending. When CBA filed their product paperwork on M-Shwari with CBK, the loan fee was called a facilitation fee and not an interest rate so their 7.5% monthly fee was allowed to stand. For the other banks' digital products, they called the charge an interest rate when they applied for CBK approval and were therefore subject to the regulation. On the savings side, due to the regulation, interest rates were raised to 6.2% and all banks implemented this across all their savings products, including across all digital platforms. M-Shwari was, therefore, subject to the floor on deposit rates, just like its direct competitors in the digital platform space.

Since M-Shwari did not change its loan facilitation fee as a result of the regulation, it did not change its screening of customers. However, as expected, other banks like KCB and Equity, adjusted their digital lending practices once the interest rates on these loans fell by almost an order of magnitude. They restricted their digital lending to focus on lower risk individuals (this was shared with us in conversations with a lender). Given all these digital platforms use the same underlying data (from M-PESA) to build credit scores, the change due to the interest rate caps was simply to raise the credit score cutoff at which new loans were approved. There could also have been accompanying changes in the loan limits going forward. Finally, all these digital platforms implemented identical deposit rates for savings, including M-Shwari. The only distinguishing feature of M-Shwari from its direct competitors was the interest it charged on loans (7.5% per month vs. close to 1% a month), but not the interest it paid out on deposits. M-Shwari therefore had a policy "carve out" when it came to lending rates, but not for deposit rates. In this paper, we use this natural experiment of the carve out for M-Shwari, along with a theoretical model that we estimate structurally, to understand the role of interest rate caps as well as the welfare consequences of this carve out.

3 Data

We use two separate de-identified administrative datasets on M-Shwari accounts, both provided by CBA, at the account level, for random subsets of customers. We have no comparable administrative data from other banks.

The first dataset covers a random sample of almost 10,000 customers who opened accounts between January and March 2016, with complete loan histories and credit limits. From this, we construct a customer-day panel tracking daily loan and credit limit evolution from January 3, 2016 to March 31, 2017. The empirical analysis uses a regression discontinuity design around the implementation date of the interest rate caps at the day level. Table 1 reports summary statistics for this credit data within the RD window (46 days on either side of implementation). On any given day in this period, about 2% of customers had a loan; the average loan amount was KShs 20 (including zeros), and the average loan limit was KShs 836.

The second dataset covers a random sample of 5,000 customers from the same account-opening period, with daily savings balances. We cannot match savings and credit accounts, as the bank restricted this linking to prevent reconstruction of its credit scoring rules. We construct a customer-day panel for savings, with Table 1 showing RD-window summary statistics; the average daily savings balance was KShs 624.

Because all data are daily and the RD bandwidth is 46 days, our estimates capture short-run responses to the caps. We cannot assess longer-term effects such as entry or exit from the banking system. However, the main digital lending products available during the regulation period remain in the market today and continue to dominate digital lending.

We also use survey data collected in Suri, Bharadwaj, and Jack (2021) to complement the administrative data for the structural estimation. From the survey, we take average loan outcomes (number of loans and average loan size) by credit score, separately for M-Shwari, its main competitors (KCB and Equity), and other lenders (e.g., friends and family, savings and credit cooperatives, microfinance institutions). We also use aggregated values of the variables M-Shwari employs to compute credit scores, by credit score. Below we describe in more detail how these data are used in the structural estimation.

4 Reduced Form Effects of the Interest Rate Caps

The first step of our empirical analysis is to document the reduced form impacts of the interest rate caps on financial behavior, using the administrative data described above. To understand the impacts of the interest rate caps, we use an RD design around the date the caps were implemented. While the caps were announced three weeks before before their implementation, there was some uncertainty around the exact implementation and we find no evidence of anticipatory effects (results available upon request).

We use these reduced form effects to establish some intuition for the effects we expect to find. In the next section, we build a formal model to rationalize these results. Although we only have administrative data from M-Shwari, since it has the largest market share at the time, and given the largest competitors in digital lending are using similar non-traditional data to score their customers (in the case of KCB, they are using exactly the same data from Safaricom), studying M-Shwari tells us not just the impacts of the interest rate caps on M-Shwari customers directly, but also the competitive responses that M-Shwari may undertake in response.

In terms of consumer responses, with the interest rate caps, we expect that M-Shwari's competitors will raise their credit score cutoff, making it harder for consumers to borrow from them. We therefore expect these consumers to try to borrow from M-Shwari, with M-Shwari therefore making more and larger loans post the interest rate caps. If we think of those who have higher credit scores, we expect those consumers to leave M-Shwari as they may still be eligible for loans from M-Shwari's competitors, and these loans are now much cheaper (1.1% vs. 7.5%). In terms of savings behavior, we expect that lower credit score consumers will save more in M-Shwari to try and build up their credit limits since they are less likely to be able to borrow from M-Shwari's competitors. Note that the savings behavior would not respond differentially across lenders purely due to the regulated floors on deposit rates as these were implemented across both M-Shwari as well as its competitors.

We therefore begin by examining individuals' financial behavior before and after the implementation date. Hence, we estimate the following equation:

$$Y_{it} = \alpha + \beta D_t + \gamma_1 (M_t - d) + \gamma_2 (M_t - d) * D_t + \epsilon_{it} \tag{1}$$

where Y_{it} is an outcome for individual i at time t, and D_t is an indicator variable for whether the date of observed activity in the M-Shwari account is after the implementation of the interest rate caps, d, and M_t are standard ways for controlling for trends in borrowing behavior in the days leading up to and after the regulation. So, γ_1 and γ_2 flexibly capture the direct effect of the "running variable" (in this instance calendar days) on the outcome of interest. Given these controls, β captures the "post interest rate cap" effect, or the treatment effect of interest rate caps. In the results, we only report the β coefficient. This regression is the standard local linear specification commonly used in RD designs and our regressions show results for the optimal bandwidth (46 days on either side of the cutoff).

In order to examine heterogeneity by the risk of the customer, we also estimate versions of Equation 1 above where we interact the D_t with X_i where X_i denotes measures of risk, like the individual's original credit score.

$$Y_{it} = \alpha + \phi D_t * X_i + \psi D_t + \omega X_i + \gamma_1 (M_t - d) + \gamma_2 (M_t - d) * D_t + \epsilon_{it} \tag{2}$$

In this instance, the coefficient of interest is ϕ which captures the differential effect of those with higher credit scores (say) *after* the regulation is implemented. For all our specification, we report standard errors clustered by individual.

4.1 Reduced Form Results

Figure 1 shows our main result graphically. Each dot represents a day and the vertical line the date of implementation (September 14 2016). The top panel shows that, after the implementation of the caps, the probability of having a loan with M-Shwari increases as do loan amounts (the bottom panel). Note that, using simulations, it has recently been shown that visual representation of RD effects are unclear when the t-statistics are small (in the 2-4 range).[2] Table 2 presents the regression versions of these figures.

[2] See https://twitter.com/KiraboJackson/status/1074062192037847040. These simulations were first posted on Twitter by Kirabo Jackson in 2018.

The first column in Table 2 shows the estimates of β from Equation 1, which we interpret as the overall "post interest rate cap" effect. In Column 1, this coefficient is 0.0018 indicating that the probability that an individual takes a loan from M-Shwari on a given day increases by 0.18 percentage points after the implementation of the cap. Since the baseline probability of having a loan is 1.8%, this represents a nearly 10% increase in the probability of having an M-Shwari loan. The next column interacts the interest rate cap date with an indicator variable for whether someone was eligible for a loan under M-Shwari at the time they opened an account. At the bottom of Table 2, we show the F test (and the corresponding p-value) on the overall effects of the interest rate caps for specifications where we include interactions.

As described, M-Shwari uses a threshold or cutoff in their credit score to determine whether someone is eligible for a loan. Any score above zero makes a customer eligible, though customers do not know their scores, just their approved credit limit. Column 2 of Table 2 shows that individuals above the credit score cutoff are *less* likely than those with scores below to take a loan after the interest rate caps. This is in line with the idea that those with higher scores are less likely to use M-Shwari since they are less risky borrowers and have far cheaper alternative credit options now. Column 3 estimates a similar regression, but instead of the credit score cutoff, it uses the credit score directly. As expected, the interaction term is negative and significant, indicating that those with better credit scores are less likely to take up a loan with M-Shwari after the implementation of the caps. Column 4 shows estimates for the sub sample that has positive credit scores, which are similar to those for the overall sample, reinforcing the idea that less risky borrowers seek other sources of credit, while riskier borrowers increase their take up of M-Shwari. Columns 5-8 in Table 2 replicate columns 1 through 4, but with loan amounts as the outcome. Column 5 shows that, after the implementation of the caps, loan amounts went up by nearly 1.1%. However, those with better credit scores take out less in loans.

Columns 9 through 12 in Table 2 examine the impacts on credit limits, which are set by M-Shwari and are raised or decreased depending on consumer behavior. There are many ways for a customer to increase credit limits, namely by increasing savings, and repaying loans on time. Credit limits therefore potentially give us insight into both consumer behavior as well as the supply response of M-Shwari itself as the bank ultimately makes the choice to change limits. Column 9 shows that credit limits increase after the implementation of the interest rate caps, but as Columns 10 through 12 indicate, the average increase is driven mainly by individuals who have higher credit scores. Column 11 shows that those with an additional point in credit scores see their credit limits increase by 117 KSh, a nearly 15% increase from the average limit. M-Shwari actively tried to retain low risk consumers after the implementation of the interest rate caps by offering them higher limits.

In summary, Table 2 demonstrates our headline findings: loans, loan amounts and credit limits increase in response to the interest rate caps, but those with better credit scores are taking fewer loans, taking less in loan amounts and are approved for higher limits in an attempt to keep them on the platform.[3]

Table 3 looks more carefully at the idea that individuals with high credit scores behave differently from people with lower credit scores in response to the interest rate caps. We look at how the interest rate caps affect outcomes for those across the four quartiles of the initial credit score. At the bottom of the table we report the F test (and corresponding p-value) of the overall effect for those in the top credit score quartile. Column 1 shows that, while there is an overall increase in loan take up after the caps are in place (coefficient of 0.0027), those in the highest credit score quartile are significantly *less* likely to take up credit (coefficient of -0.0019). Column 2 shows the same result but for the sample with positive credit scores. We see the same pattern for loan amounts and for credit limits. In Column 3, we find that while overall loan amounts go up by 1.6%, those in the highest quartile of credit scores reduce their loan amounts by 1.18%. There is therefore an offset here: while M-Shwari gains in terms of loans taken out and loan amounts, these appear to exclusively come from the high risk customers. Columns 5 and 6 explore credit limits. Those in the highest quartile of credit scores see their credit limits *increase* once interest rate caps are in place. We interpret this as a supply side response by the bank in an attempt to keep these customers borrowing from M-Shwari.

Next, we try to understand the strategic behavior of customers in response to the caps, starting with savings behavior. One way to improve credit limits with M-Shwari is to save more. Hence, we expect that those with the lowest credit scores save more as they have the greatest incentives to grow their credit limit with M-Shwari. Column 1 in Table 4 shows that the probability of positive savings in M-Shwari on a given day increases significantly after the implementation of the caps (the increase is 1.5 percentage points). Columns 2 and 3 explore heterogeneity by credit score, and as predicted, individuals with high credit scores save *less* in the system as they have less of an incentive to do so now because they prefer to build borrowing histories and credit limits in cheaper platforms. We find similar results for savings amounts in columns 5-8. While overall savings increase after the cap comes into place, those with worse credit scores are the ones saving more as a way to improve their credit limit.

[3] In Appendix Table A1 we show that individuals are more likely to access M-Shwari loans for the *first time* after the interest rate caps go into effect. However, first time loan takers do not seem to be just low risk types. In magnitude, individuals are 0.07 percentage points more likely to take up M-Shwari for the first time after the implementation of the caps, which represents a 70% increase.

5 Theory

Given the reduced form results above, we now build a model to rationalize these results. We then structurally estimate the model to be able to perform counterfactuals. In this context, M-Shwari was provided a carve out from the regulation, i.e. an exemption from being subject to the interest rate caps on lending that the entire industry was subject to. As a reminder, they were still subject to the same floor on interest rates on deposits as everyone in the industry. In our counterfactual analysis, we simulate out what would have happened had M-Shwari not been provided this exemption or carve out and hence estimate what the welfare impacts of the carve out are on consumers. To do this, we need to understand the competitive landscape M-Shwari operates in, and not just how M-Shwari itself responds, but also how its (representative) competitor responds to the regulations. In this section, we therefore build out a simple model of the digital credit product industry, modeling directly M-Shwari and its competitors, which then helps us analyze both demand and supply side responses to the regulation. We then estimate this model structurally.

5.1 Model Setup

There are two firms $\{1,2\}$, where we use $j=1$ to denote M-Shwari and $j=2$ to denote the competitor. There is a continuum of borrowers indexed by a measure of creditworthiness r (i.e. this is their credit score, and also what we refer to as a market below).

Each borrower wants to borrow from one of the two firms. A loan of size m gives the consumer the following utility

$$u_j = \beta m(v - i) + \xi_j + \varepsilon$$

where v is the per unit value-in-use of the loan, i is the interest rate, ξ_j are firm-specific preference parameters, β captures the sensitivity of demand and ε is an error, distributed type 1 extreme value. We assume $v > i_j$ for each firm, so that borrowers derive positive surplus from a loan.

The probability of repayment of a loan is given by the function $p(r,m)$, where p is strictly concave and decreasing in m for each r. This represents the idea that the borrowers are more likely to repay smaller loan amounts, and have a convex propensity to default as loan sizes increase. We assume p is strictly increasing in r for each m i.e. this is truly a measure of creditworthiness. We also assume p is supermodular in (r,m) i.e at higher loan amounts the borrowers become more differentiated in their risk of default.

Interest rates are set exogenously to be at i_n for firm n. This is without loss in the application we will consider, because in the period prior to the imposition of the interest rate caps, the interest rates are approximately fixed at the observed rate and we use this to back out the other margins of adjustment by the firms. In the period post interest rate caps, the interest rates are moved exogenously by policy. Thus, the instruments (or choices) available to the firms are:

1. A cutoff $\underline{r}$ below which they do not lend.
2. Credit limits $\bar{m}(r)$ at each $r \geq \underline{r}$ (else they are zero).

From the data, we see that the share of the outside option is stable across markets (credit scores) and across time. Thus, we focus on only the inside options for this analysis, normalizing their shares by the inside market share. Moreover, we keep v as fixed, even though in the estimation we will allow v to be drawn from a distribution.

Consider the borrower problem, conditional on the equilibrium loan terms chosen by the two firms. Given r, the market shares of the two products can be backed out from

$$s_j(r, m_1, m_2) = \begin{cases} \dfrac{\exp[\beta(v - i_1)m_1 + \xi_1]}{\sum_j \exp\left[\beta(v - i_j)m_j + \xi_j\right]} & \text{if } r \geq \max(\underline{r}_1, \underline{r}_2) \\ 1 & \text{if } r \in [\underline{r}_1, \underline{r}_2] \\ 0 & \text{otherwise} \end{cases}$$

where $s_j(r, m_1, m_2)$ is firm j's share of the market.

Knowing this, firm j's problem at any given r is

$$\max_{m_j} \Big[m_j \Big(i_j\, p(r, m_j) - c \Big) s_j(r, m_1, m_2) \Big].$$

Assuming the first order approach is valid, the optimal credit limit $\bar{m}_j(r)$ solves the following equation:

$$\Big[i_j\, p(r, m_j) + i_j\, m_j \frac{\partial p(r, m_j)}{\partial m_j} - c \Big] s_j(r; m_1, m_2) + \Big[i_j\, p(r, m_j) - c \Big] m_j \frac{\partial s_j(r; m_1, m_2)}{\partial m_j} = 0. \quad (3)$$

Finally, the credit score cutoff $\underline{r}_j$ solves

$$i_j\, p(r, 0) - c = 0. \quad (4)$$

5.2 Theoretical Results

We first investigate the static model's equilibrium in the period prior to the interest rate caps, before moving on to analyze the impact of the policy change itself. We then incorporate strategic behavior by the borrowers. All proofs are in the Appendix.

Static equilibrium: Driven by empirical facts, we make the following assumptions:

Assumption 1. *1. $i_1 > i_2$ where firm 1 is M-Shwari and firm 2 is the competitor*

2. $s_1(r, m_1, m_2) > \frac{1}{2}$ for all r, i.e. firm 1 is the dominant firm in the market[4]

Consider the "monopoly markets" i.e. those credit scores at which only one of the firms offers a product. These will generally be the borrowers with the lowest credit scores. In these markets, we have the following result:

Proposition 1. *Monopoly markets are served by firm 1.*

Effectively, firm 1 can afford to serve riskier customers because it is charging a higher interest rate. Firm 1's credit limits also follow intuitive patterns:

Proposition 2. *The credit limits $\bar{m}_1(r)$ are increasing in r.*

Policy impact: We now turn to analyzing the impact of the interest rate policy change. Suppose there is an exogenous shift that moves i_2 to $i_2' < i_2$ while leaving $i_1' = i_1$.

Proposition 3. *Provided i_2' is high enough that firm 2 remains active (i.e. $i_2' p(r, 0) \geq c$ for some r), $\underline{r}_2' > \underline{r}_2$, i.e. firm 2 tightens up its lending cutoff, while $\underline{r}_1' = \underline{r}_1$.*

As expected, the second firm must now tighten its lending practices in order to be able to lend at a lower rate. The first firm is unaffected at the margin. The two propositions above map directly to our reduced form results above that individuals with lower credit scores are more likely to borrow and borrow more from M-Shwari (firm 1) after the interest rate caps go into effect (Table 2 and Table 3).

However, there are interesting effects in the markets that firm 2 is forced to exit. The next proposition shows that all borrowers who are now in monopoly markets become strictly worse off. Not only does firm 2 no longer provide its product in these markets, but firm 1 also decreases its credit limits in these markets. Intuitively, competition in these markets previously enabled borrowers to have more options and better deals from each option. Thus, although a carve out for firm 1 to continue charging higher interest rates ensures that riskier borrowers maintain access to credit, that credit now comes at strictly worse terms through the channel of market competition. This explains some of the reduced form results in the last columns of Table 2 and Table 3.

Proposition 4. *Consider $r \in (\underline{r}_2, \underline{r}_2')$. The following hold for such borrowers:*

[4] This is effectively an assumption on i and ξ.

1. $\bar{m}_2'(r) = 0 < \bar{m}_2(r)$*, i.e. their cap with firm 2 strictly decreases (to* 0*).*

2. $\bar{m}_1'(r) < \bar{m}_1(r)$*, i.e. their cap with firm 1 strictly decreases.*

We now analyze what happens to borrowers in the markets where both firms continue to operate, which is more complex. A decrease in firm 2's interest rate has competing effects on firm 1's credit limit-setting. On the one hand, firm 2 becomes a more attractive competitor, which might incentivize firm 1 to increase its credit limits to maintain market share. On the other hand, firm 2 will itself reduce its credit limits in response to its lower interest rate — since lower revenue per loan necessitates tighter lending — and since credit limit-setting is a game of strategic complements, firm 1's best response moves in the same direction. The net effect on firm 1's credit limits, and hence on borrower welfare, is therefore an empirical question. Nevertheless, this analysis demonstrates that, in the best-case scenario, the gains accruing to richer, less risky borrowers from a lower interest rate on firm 2's loans could mitigate the adverse effects on poorer, riskier borrowers, whereas in the worst case, all borrowers could be worse off.

Borrower response: Finally, suppose borrowers can privately signal higher creditworthiness to each firm by saving more with that firm, increasing their perceived credit score from r to $r + a_j$ at a strictly convex cost $c(a_j)$. This fits the institutional context, where credit scores are firm-specific and not publicly observable. Concurrently with the interest rate caps, the regulation introduced deposit rate floors, which we model as a decrease in the marginal cost of saving. Given optimal credit limits and firm choice, the borrower solves:

$$\max_{a_{j,t}} \beta\, \bar{m}_{j,t}(r + a_{j,t})(v - i_j) - c_{j,t}(a_{j,t}).$$

We assume that $c_{j,t}$ is sufficiently convex such that this objective is strictly concave in $a_{j,t}$, ensuring a unique interior solution.

Proposition 5. *Consider a borrower with* $r \in (\underline{r}_2, \underline{r}_2')$ *who borrows from firm 1. If the reduction in the marginal cost of saving strictly exceeds any decrease in the marginal benefit,*

$$\Delta c'(a^*_{1,\text{pre}}) \;>\; \beta \left[\frac{\partial \bar{m}_{1,\text{pre}}}{\partial r}(r + a^*_{1,\text{pre}}) - \frac{\partial \bar{m}_{1,\text{post}}}{\partial r}(r + a^*_{1,\text{pre}}) \right]^+ (v - i_1),$$

where $\Delta c'(a) \equiv c'_{1,\text{pre}}(a) - c'_{1,\text{post}}(a) > 0$ *and* $[x]^+ = \max\{x, 0\}$*, then* $a^*_{1,\text{post}} > a^*_{1,\text{pre}}$.

The condition is satisfied whenever the marginal credit-limit response $\frac{\partial \bar{m}_1}{\partial r}$ does not decrease from the duopoly to the monopoly regime, and more generally requires only that the deposit rate floor provides a large enough cost reduction to compensate for any weakening of competitive incentives. This matches the reduced form results in Table 4.

6 Structural Estimation

In this section, we develop a rigorous empirical framework to recover the latent utility parameters that govern borrowers' choices in the presence of interest rate caps and lending limits, which will then allow us to speak to important welfare effects and counterfactual scenarios. Our approach builds on a two-step Generalized Method of Moments (GMM) estimation strategy, which integrates simulation-based inversion of market shares with moment condition construction. The methodology exploits both cross-sectional heterogeneity (across credit scores, or markets in the theory) and time-series variation (pre- and post-policy change) in the data for identification.

6.1 Setup and Assumptions

The utility derived by a borrower n from a loan of size m offered by firm $j \in \{1,2\}$ in market r is given by

$$u_{n,j,r} = \beta \bar{m}_{j,r}(v_{n,j,r} - i_j) + \xi_{j,r} + \varepsilon_{i,j,r} \tag{5}$$

where β is a sensitivity parameter that governs how loan size m_j interacts with the deviation of the latent value v from its mean μ; $\bar{m}_{j,r}$ are the lending limits chosen by the firms in each market; v is a latent "value-in-use" variable, modeled as below for each market r. ; $\xi_{j,r}$ are product-market specific unobservables, and $\varepsilon_{n,j,r}$ is a logit shock.

$$v_{n,j,r} = \mu + \sigma \nu, \quad \nu \sim \mathcal{N}(0,1), \tag{6}$$

where μ represents the central tendency of v, and $\sigma > 0$ captures the dispersion.

The model thus reflects that borrowers' utility is driven not only by observed loan characteristics (such as the credit limit and the interest rate) but also by unobserved demand parameters that can influence the credit limits m that are chosen by the firms. The utility from the outside option is normalized to 0. The observed default rates in each market will function as cost-side instruments that move around the credit limits, m, chosen by the firms.

6.2 Data for the Structural Estimation

Our empirical analysis exploits three complementary data sources that, in concert, provide a comprehensive view of the lending environment and inform our estimation strategy.

First, the *survey data* described above captures detailed information on borrowers' creditworthiness and lending activity. For a range of credit scores r, we can compute the number of loans provided by the different lenders. Specifically, we denote by $L_1(r)$ the number of loans from firm 1 (M-Shwari) in credit score r, by $L_2(r)$ the number of loans from firm 2 (the competitor), and by $L_0(r)$ the number of loans from alternative (outside) lenders. In addition, the survey reports the corresponding average loan sizes, $\bar{m}_1(r)$ for the incumbent, $\bar{m}_2(r)$ for the competitor, and $\bar{m}_0(r)$ for the outside options. We then construct the market shares for the inside options as follows:

$$s_1(r) = \frac{L_1(r)}{L_1(r) + L_2(r) + L_0(r)}, \quad s_2(r) = \frac{L_2(r)}{L_1(r) + L_2(r) + L_0(r)}. \tag{7}$$

Second, the *administrative data* provide precise measures of loan performance. In particular, for each credit score r, we observe the variables used by the firm to determine that credit score, which we collect into an instrument vector z_r. These include M-PESA transaction volumes and values, paybill counts and amounts, bank transfer activity, and send/receive patterns. These supply-side variables serve as instruments in our demand-side moment conditions, as we describe below.

Third, the *policy change data* record key changes in the lending environment surrounding the interest rate caps. For each credit score r, the administrative data includes the number of loans before and after the policy change for firm 1, which we denote by $N_{\text{pre}}(r)$ and $N_{\text{post}}(r)$, respectively. From these, we compute the proportional change in loan counts as

$$\Delta s_1(r) = \frac{N_{\text{post}}(r) - N_{\text{pre}}(r)}{N_{\text{pre}}(r)}. \tag{8}$$

and set $\Delta s_1(r) = -\Delta s_2(r)$, since almost all the substitution during the policy intervention happens across the inside options. Moreover, we observe that the shares for firm 1 increase below score $\hat{r} = 0$ and decrease above $\hat{r}$, so we infer that firm 2 exited from markets $r < \hat{r}$ after the policy intervention, due to the decreased reward from staying in those markets not being worth the risk.

These three data sources are merged using the credit score r as the common identifier, thereby forming the basis of a two-period panel. In the pre-policy period, the markets are characterized by the survey and administrative measures:

Credit quality: r,

Market shares: $s_1(r)$, $s_2(r)$,

Default rate: $p(r)$,

Credit limit firm 1: $\bar{m}_1(r)$,

Credit limit firm 2: $\bar{m}_2(r)$.

The interest rates are set exogenously at $i_1 = 0.075$ for the incumbent and $i_2 = 0.07$ for the competitor. In the post-policy period, the interest rate for the incumbent remains unchanged at $i_1 = 0.075$, whereas the competitor's rate is reduced to $i_2 = 0.012$ to reflect the policy intervention. The market shares are updated to $s_1^*(r)$ and $s_2^*(r)$ using $\Delta s_i(r)$, while the other variables (e.g., default rates, credit limits) remain consistent with their pre-policy values or are directly observed. This comprehensive panel structure—exploiting both cross-sectional variation across credit scores and temporal variation induced by the policy—forms the empirical foundation for our subsequent estimation of the latent utility parameters.

6.3 Market Share Inversion and Moment Construction

We now discuss the details of how we estimate the structural model. We define

$$\delta_{j,k} = \beta \bar{m}_{j,k}(\mu - i_j) + \xi_{j,k}$$

i.e. the mean utility to product j in market k. A central challenge in our estimation is to recover the unobserved $\delta_{j,k}$ parameters from observed market shares. Given Equation (5), the predicted market share for firm j is modeled via a multinomial logit:

$$s_j = \frac{\exp(u_j)}{1 + \exp(u_1) + \exp(u_2)}, \tag{9}$$

where the denominator accounts for an outside option with utility normalized to zero. In practice, we approximate the expectation over the latent variable v by drawing $N = 1000$ normal variates. For each draw, the utility u_j is computed and the corresponding share s_j is evaluated.

To invert the observed shares s_j^{obs}, we solve the following minimization problem:

$$\min_{\delta} \left\{ \left(s_1^{\text{obs}} - \frac{1}{N} \sum_{k=1}^{N} \frac{\exp(\delta_1 + \beta\, m_1\, (v_k - \mu))}{D_k} \right)^2 + \left(s_2^{\text{obs}} - \frac{1}{N} \sum_{k=1}^{N} \frac{\exp(\delta_2 + \beta\, m_2\, (v_k - \mu))}{D_k} \right)^2 \right\}, \tag{10}$$

where $D_k = 1+\exp(\delta_1+\beta\, m_1\,(v_k-\mu))+\exp(\delta_2+\beta\, m_2\,(v_k-\mu))$. This inversion yields estimates $\hat{\delta}_1$ and $\hat{\delta}_2$ for each market in both periods.

Subsequently, we construct moment conditions that exploit both cross-sectional and time-series variation:

1. **Time-Difference Moment:** For competitive markets where firm 2 remains active after the policy change ($r > \hat{r}$), and where the credit limit $\bar{m}_2$ does not change across periods, differencing the mean utility equation across periods yields

$$\Delta\delta_2 + \beta\, \bar{m}_2\, \Delta i_2 = 0, \tag{11}$$

   where $\Delta\delta_2 \equiv \delta_{2,r,\text{post}} - \delta_{2,r,\text{pre}}$ is the change in firm 2's recovered mean utility, and $\Delta i_2 \equiv i_{2,\text{post}} - i_{2,\text{pre}}$ is the exogenous policy-induced change in firm 2's interest rate. This condition follows from the assumption that $\xi_{j,r}$ is time-invariant (see Section 6.4).

2. **Instrumental Variable Moments:** For each firm $j \in \{1,2\}$, we impose

$$\mathbb{E}\left[z_r \cdot \xi_{j,r}\right] = \mathbb{E}\left[z_r \cdot \left(\delta_{j,r} - \beta\, \bar{m}_{j,r}\,(\mu - i_j)\right)\right] = 0, \tag{12}$$

   where z_r is the vector of supply-side instruments described in Section 6.2. The exclusion restriction is that these variables—which the firm uses to determine credit scores and hence credit limits—affect loan terms through the supply side but do not directly enter borrower utility. The relevance condition—that these variables predict credit limits—follows from the firm's credit-limit-setting rule in Equation (3).

Together, the time-difference condition and the instrumental variable moments, evaluated across all markets and both time periods, form the basis of our GMM estimation.

6.4 Identification and Estimation of Model Parameters

Before describing the estimation procedure, we discuss the identification of the structural parameters $\theta = (\beta, \mu, \sigma)$. The argument rests on three assumptions, which we state here informally; formal statements and proofs are provided in Appendix 7.

First, the demand unobservable $\xi_{j,r}$ is *time-invariant*: unobserved product-market characteristics do not change between the pre- and post-policy periods. This is plausible given the short time horizon spanned by the regression-discontinuity window around the policy change, and is a standard assumption in panel demand models.

Second, the instruments z_r satisfy the *exclusion restriction* $\mathbb{E}[z_r \cdot \xi_{j,r}] = 0$. These supply-side variables—M-PESA transaction patterns, paybill activity, and similar measures used by the firm to construct credit scores—affect credit limits through the firm's profit-maximizing screening rule (Equation 3) but do not directly enter borrower demand. Conditional on the loan terms $(\bar{m}_j, i_j)$ a borrower faces, the underlying transaction data used to assign the credit score does not affect which firm the borrower prefers.

Third, the instruments are *relevant*: $\mathbb{E}[z_r \cdot \bar{m}_{j,r}] \neq 0$. This follows from the model, since the credit limit $\bar{m}_{j,r}$ is a function of the credit score r, which is itself determined by z_r.

Sources of identification. Consider first the share inversion (Equation 10). For any candidate values of (β, σ), the observed shares and loan characteristics uniquely determine the mean utilities $\delta_{j,r,t}$. The random coefficient in the share formula depends on β and σ only through the product $\beta\sigma$ (via the term $\beta\bar{m}_j\sigma\nu$), so we can write the recovered mean utilities as functions of this composite parameter: $\delta_{j,r,t} = \delta_{j,r,t}(\beta\sigma)$.

The *sensitivity parameter* β is identified primarily through the exogenous temporal variation induced by the policy change. For firm 2 in competitive markets, the time-difference condition (Equation 11) gives:

$$\delta_{2,r,\text{post}}(\beta\sigma) - \delta_{2,r,\text{pre}}(\beta\sigma) = -\beta\,\bar{m}_{2,r}\,\Delta i_2.$$

The left-hand side depends on $\beta\sigma$ through the share inversion; the right-hand side depends on β alone (and observables). This equation separates β from the composite $\beta\sigma$. In the limiting case $\sigma \to 0$, the share inversion reduces to the standard logit formula $\delta_j = \log(s_j/s_0)$, which is independent of any parameters, and the time-difference directly determines β from observable share changes and the exogenous interest rate shift. More generally, requiring this condition to hold simultaneously across competitive markets with different credit limits $\bar{m}_{2,r}$ pins down both β and σ.

The *mean valuation* μ is identified from the cross-sectional instrumental variable moments (Equation 12). Given (β, σ), the mean utilities $\delta_{j,r}$ are known, and $\xi_{j,r} = \delta_{j,r} - \beta\bar{m}_{j,r}(\mu - i_j)$. The orthogonality condition $\mathbb{E}[z_r \cdot \xi_{j,r}] = 0$ then determines μ from the instrumented cross-market relationship between $\delta_{j,r}$ and $\bar{m}_{j,r}$. Intuitively, μ governs the slope of $\delta_{j,r}$ with respect to the credit limit (since $\delta = \beta\bar{m}(\mu - i) + \xi$), and the instruments provide exogenous cross-market variation in $\bar{m}$ to estimate this slope.

The *dispersion parameter* σ governs the extent to which the model departs from the Independence of Irrelevant Alternatives (IIA) property of the standard logit. When $\sigma > 0$, the random coefficient $\beta\bar{m}_j\sigma\nu$ creates substitution patterns that depend on the relative loan sizes of the two firms: borrowers with high ν draws disproportionately favor the firm offering larger loans. Markets with different $(\bar{m}_1, \bar{m}_2)$ vectors therefore exhibit different substitution patterns under $\sigma > 0$ but identical (IIA) patterns under $\sigma = 0$. This cross-market variation in the sensitivity of shares to loan characteristics identifies σ. With only two inside options per market, this source of identification is inherently weaker than in settings with many products, which is reflected in the larger standard error on $\hat{\sigma}$.[5]

Joint identification. While the discussion above attributes each parameter to a primary source of variation, identification is ultimately joint: all three parameters enter the moment conditions simultaneously through the share inversion. Appendix 7 establishes formally that, under the assumptions stated above and a support condition on the cross-market variation in credit limits, the moment conditions generically have a unique solution in (β, μ, σ).

We estimate the parameter vector $\theta = (\beta, \mu, \sigma)$ using the following two-step GMM procedure:

Step 1. The initial estimator $\theta^{(1)}$ is obtained by minimizing the quadratic form

$$Q_1(\theta) = \bar{m}(\theta)'\bar{m}(\theta), \tag{13}$$

where $\bar{m}(\theta)$ denotes the sample average of the moment conditions, and the weighting matrix is the identity matrix.

Step 2. Using the residuals from Step 1, we compute an estimate of the covariance matrix of the moment conditions, denoted by S. The optimal weighting matrix is then given by $W = S^{-1}$. The final estimator $\hat{\theta}$ is obtained by minimizing

$$Q_2(\theta) = \bar{m}(\theta)' W \, \bar{m}(\theta). \tag{14}$$

We compute standard errors for $\hat{\theta}$ via the sandwich estimator:

$$\mathrm{Var}(\hat{\theta}) = (G'WG)^{-1} (G'WSWG) (G'WG)^{-1}, \tag{15}$$

where G is the Jacobian matrix of the moment conditions with respect to θ, approximated using finite differences.

6.5 Results and Counterfactual

We first estimate the main parameters in the model, specifically β, μ and σ from above, along with their accompanying standard errors, reported in Table 5. The estimated mean of the latent value, μ is 0.248, while the dispersion parameter, σ, is substantially smaller, at 0.01, suggesting

[5]More precise identification of σ could be achieved with finer variation, for example at the province-level; we show below that our welfare conclusions are robust to σ across its plausible range.

that there is low heterogeneity in borrowers' unobserved valuation of loans. The mean value being high matches our observation that outside option shares are not very responsive to credit limits and interest rate changes: the people who do end up taking loans from the two firms have a very high value for them relative to the outside option. The estimate of β is 0.001, implying that at a representative loan size of 3000 Kenyan shillings in the dataset, a reduction in the interest rate by 0.01 leads to an increase in latent utility by 0.03. The estimates for β and μ are obtained with reasonable precision, but the standard error for $\hat{\sigma}$ is large. As discussed in Section 6.4, σ is identified from cross-market variation in how shares respond to different loan-size vectors, which is a relatively weak source of variation with only two inside options per market.

Our primary motivation in performing the structural estimation is to be able to analyze the unique policy decision to exempt one firm from the interest rate caps, i.e. the M-Shwari carve out. In particular, we can now counterfactually predict what would have happened had there been a uniform interest rate cap, and then compare the ways in which this carve out improves or worsens outcomes for borrowers, and for which types of borrowers.

We simulate the uniform cap counterfactual by capping both firms interest rates at the second firm's post-policy interest rate. As we can see in Equation 4, the cutoff credit score below which a firm does not lend is determined only by the interest rate, so we can use the second firm's post-period cutoff score as the first firm's cutoff score as well. We simulate consumer values from a Normal distribution $N(\mu, \sigma^2)$ using the estimated value parameters, and then calculate market shares in the markets that the firms continue to operate in. Borrower welfare is quantified by the expected log-inclusive value, $W = E_{v_k}[\ln(1+\sum_j e^{u_{jk}(v_k)})]$, measured in utils, where the expectation is taken over the distribution of value v_k. We weight each credit score market equally when calculating overall welfare.

We compare the following three scenarios, for not just the average borrower, but also separately for the average low credit score borrower (here, the low credit score group is defined as $r < \hat{r} = 0$; recall that these were the markets that firm 2 exited after the interest rate caps were enforced) and the average high credit score borrower:

1. The "Pre (baseline)" scenario which reflects market conditions before any policy intervention, with product mean utilities ξ_j calibrated to pre-policy market shares.

2. The "Post (actual)" scenario which captures the observed market after policy changes, including a rate cap for Firm 2 and its exit from low-credit-score markets, with ξ_j recalibrated to post-policy shares.

3. The "Post (uniform cap)" counterfactual scenario which assumes that, starting from the post-policy ξ_j values, both firms are subjected to an identical low interest rate ($i_j =$

0.012) and both are assumed to exit markets serving low-credit-score borrowers.

The welfare calculations for each of these scenarios are presented in Table 6. In the top panel we show overall welfare for the average borrower, as well as for the average low and high credit score borrower across these scenarios. In the bottom panel we show the market shares for the two firms across low and high credit score borrowers in each of these scenarios.

As we can see from the top panel, there is a small decrease in the average borrower welfare from the pre-policy baseline to the actual post-policy situation. This is driven by low-score borrowers seeing a decrease in their welfare that is actually substantial (column 2). This is mainly as a result of the low credit score borrowers losing access to the option of taking a loan from M-Shwari's competitors. The high credit score borrowers see a small improvement, since one of their options now effectively becomes cheaper due to the interest rate cap.

Looking at the counterfactual where we impose a uniform cap across all lenders, however, leads to a large reduction in overall average welfare compared to both the pre-policy baseline as well as to the actual post-policy outcome. This is primarily because the welfare for low credit score borrowers drops to zero due to both firms exiting their market. High credit score borrowers, in contrast, experience a reasonable welfare improvement under this uniform cap, since both of their options are now forced to lend at a lower rate.

Looking at the market shares reported in the bottom panel of Table 6, we see that the market shares for M-Shwari are slightly higher for the low credit score sample and zero for the competitor at baseline (the competitor do not serve this market post-policy change). However, post the policy change, for the higher credit score sample, M-Shwari has lower market shares, while the competitor has higher market shares, as they pick up some of the low-risk borrowers as their (capped) interest rates are lower. However, in the counterfactual where both M-Shwari and the competitor all have capped interest rates, neither serves the high-risk borrowers, while M-Shwari picks up a higher share of the low-risk borrowers, relative to the baseline and the market share of the competitor remains unchanged.

In summary, the structural model highlights the welfare consequences of what the reduced form results find. In particular, the carve out for M-Shwari is better for the average borrower, and more so for the high-risk (low credit score) borrowers, than if the policy had been applied uniformly across all lenders. This is driven by the core result that high-risk borrowers would be entirely excluded from the market if the policy had been uniformly applied. The carve out allows these borrowers to access credit, albeit at a higher interest rate than the rest of the market, but this ultimately benefits them. In addition, the low-risk (high credit score) borrowers are not that much worse off under the carve out, though they would be better off under a uniform cap policy as all providers would be required to offer low interest loans. Finally, in spite of M-Shwari trying to raise credit limits to keep hold of some of their low-risk borrowers, they do still lose some market share to the competitors due to the much lower interest rate that the competitors can offer these borrowers.

Because $\hat{\sigma}$ is imprecisely estimated, we verify that the welfare conclusions are robust to it across its plausible range. Table 7 reports welfare for each scenario as σ varies from 0 (standard logit) to 0.50, holding β and μ fixed at their estimated values. The welfare ranking across scenarios is stable throughout: the carve-out dominates the uniform cap for the average borrower and for low-score borrowers at every value of σ, while the uniform cap is preferred by high-score borrowers regardless of σ. Low-score borrowers are excluded under the uniform cap for all σ. The quantitative magnitudes shift—overall welfare under the carve-out ranges from 0.97 at $\sigma = 0$ to 1.29 at $\sigma = 0.50$, as higher σ adds option value—but the policy conclusions are unchanged. The welfare results are therefore driven by the interest rate shift and market exit margins rather than by the substitution patterns governed by σ.

7 Conclusion

Kenya's 2016 interest-rate regulation created an unusual policy environment: a binding cap on bank lending and a uniform deposit floor, alongside a lending-side carve-out for the dominant digital platform, M-Shwari. Using borrower-level administrative panels and an RD around the implementation date for this platform, we document that lending on the exempt platform, M-Shwari, rose (about 10% higher daily loan incidence and larger amounts), first-time access expanded (an approximately 70% increase), and credit limits were strategically adjusted—raised for safer borrowers—while lower-score customers increased savings to rebuild limits. A simple model of screening and limit-setting, estimated with administrative and survey moments, maps these reallocations into welfare: the carve-out preserved access for high-risk borrowers but produced a modest aggregate welfare decline relative to pre-policy. A uniform cap counterfactual would have generated substantially larger losses by eliminating credit entirely for low credit score borrowers.

Taken together, the evidence highlights that the incidence of rate ceilings depends not only on prices but on screening thresholds, platform competition, the ability of borrowers to take actions (saving) that relax firm-specific limits, and the supply response by the firms in the market. In markets with widespread digital access, a narrow exemption can partially offset access losses from binding caps, but at the cost of worse terms for the very borrowers who remain served. Policy designs that target screening margins directly—e.g., calibrated risk-based limits or subsidies tied to first-time formal borrowing—may better protect high-risk households than uniform price ceilings alone.

References

Alper, Emre et al. (2020). "Do Interest Rate Controls Work? Evidence from Kenya". In: *Review of Development Economics* 24.3, pp. 910–926.

Benmelech, Efraim and Tobias J. Moskowitz (2010). "The Political Economy of Financial Regulation: Evidence from U.S. State Usury Laws in the 19th Century". In: *Journal of Finance* 65.3, pp. 1029–1073.

Benston, George J. and George G. Kaufman (1996). "The Appropriate Role of Bank Regulation". In: *The Economic Journal* 106.436, pp. 688–697.

Bolen, J. Brandon, Gregory Elliehausen, and Thomas Miller (2023). "Credit For Me But Not For Thee: The Effects of the Illinois Rate Cap". Working Paper. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4315919.

Burga, Carlos, Rafael Nivin, and Diego Yamunaqué (2022). *Lending Rate Caps and Credit Reallocation*. Working Paper. Central Bank of Peru.

Campbell, John Y. et al. (2011). "Consumer Financial Protection". In: *Journal of Economic Perspectives* 25.1, pp. 91–114.

Cuesta, José Ignacio and Alberto Sepulveda (2021). *Price Regulation in Credit Markets: A Trade-off Between Consumer Protection and Credit Access*. Working Paper 21-047. SIEPR. URL: https://siepr.stanford.edu/sites/default/files/publications/21-047.pdf.

Dasgupta, Kabir and Brenden J. Mason (2020). "The Effect of Interest Rate Caps on Bankruptcy: Synthetic Control Evidence". In: *Journal of Banking and Finance* 119.

Fekrazad, Amir (2020). "Impacts of Interest Rate Caps on the Payday Loan Market: Evidence from Rhode Island". In: *Journal of Banking and Finance* 113.

Ferrari, Auror, Oliver Masetti, and Jiemin Ren (2018). *Interest Rate Caps: The Theory and the Practice*. Policy Research Working Paper 8398. The World Bank.

Grandolini, Glora M. (2015). *Five Challenges Prevent Financial Access for People in Developing Countries*. World Bank Blog.

Kuroishi, Yusuke, Cameron LaPoint, and Yuhei Miyauchi (2025). "Interest Rate Caps, Corporate Lending, and Bank Market Power: Evidence from Bangladesh". Working Paper.

Levine, Ross and Sara Zervos (1998). "Stock Markets, Banks, and Economic Growth". In: *American Economic Review*, pp. 537–558.

Madeira, Carlos (2019). "The Impact of Interest Rate Ceilings on Households' Credit Access: Evidence from a 2013 Chilean Legislation". In: *Journal of Banking and Finance* 106, pp. 166–179.

Maimbo, Samuel Munzele and Claudia Alejandra Henriquez Gallegos (2014). *Interest Rate Caps around the World: Still Popular, but a Blunt Instrument*. Policy Research Working Paper. The World Bank.

Melzer, Brian and Aaron Schroeder (2017). *Loan Contracting in the Presence of Usury Limits: Evidence from Automobile Lending*. Office of Research Working Paper 2017-02. Consumer Financial Protection Bureau.

Rigbi, Oren (2013). "The Effects of Usury Laws: Evidence from the Online Loan Market". In: *The Review of Economics and Statistics* 95.4, pp. 1238–1248.

Safavian, Mehnaz and Bilal Zia (2018). *The Impact of Interest Rate Caps on the Financial Sector: Evidence from Commercial Banks in Kenya*. Policy Research Working Paper 8393. The World Bank.

Schmukler, S et al. (2022). "Winners and Losers from Interest Rate Ceilings: Quasi-experimental Evidence from Chile". Working Paper. (2019).

Suri, Tavneet, Prashant Bharadwaj, and William Jack (2021). "Fintech and Household Resilience to Shocks: Evidence from Digital Loans in Kenya". In: *Journal of Development Economics* 153.

Temin, Peter and Hans-Joachim Voth (2008). "Interest Rate Restrictions in a Natural Experiment: Loan Allocation and the Change in the Usury Laws in 1714". In: *The Economic Journal* 118.528, pp. 743–758.

Zinman, Jonathan (2010). "Restricting Consumer Credit Access: Household Survey Evidence on Effects around the Oregon Rate Cap". In: *Journal of Banking and Finance* 34.3, pp. 546–556.

Figure 1. Loans, by Credit Score Cutoff

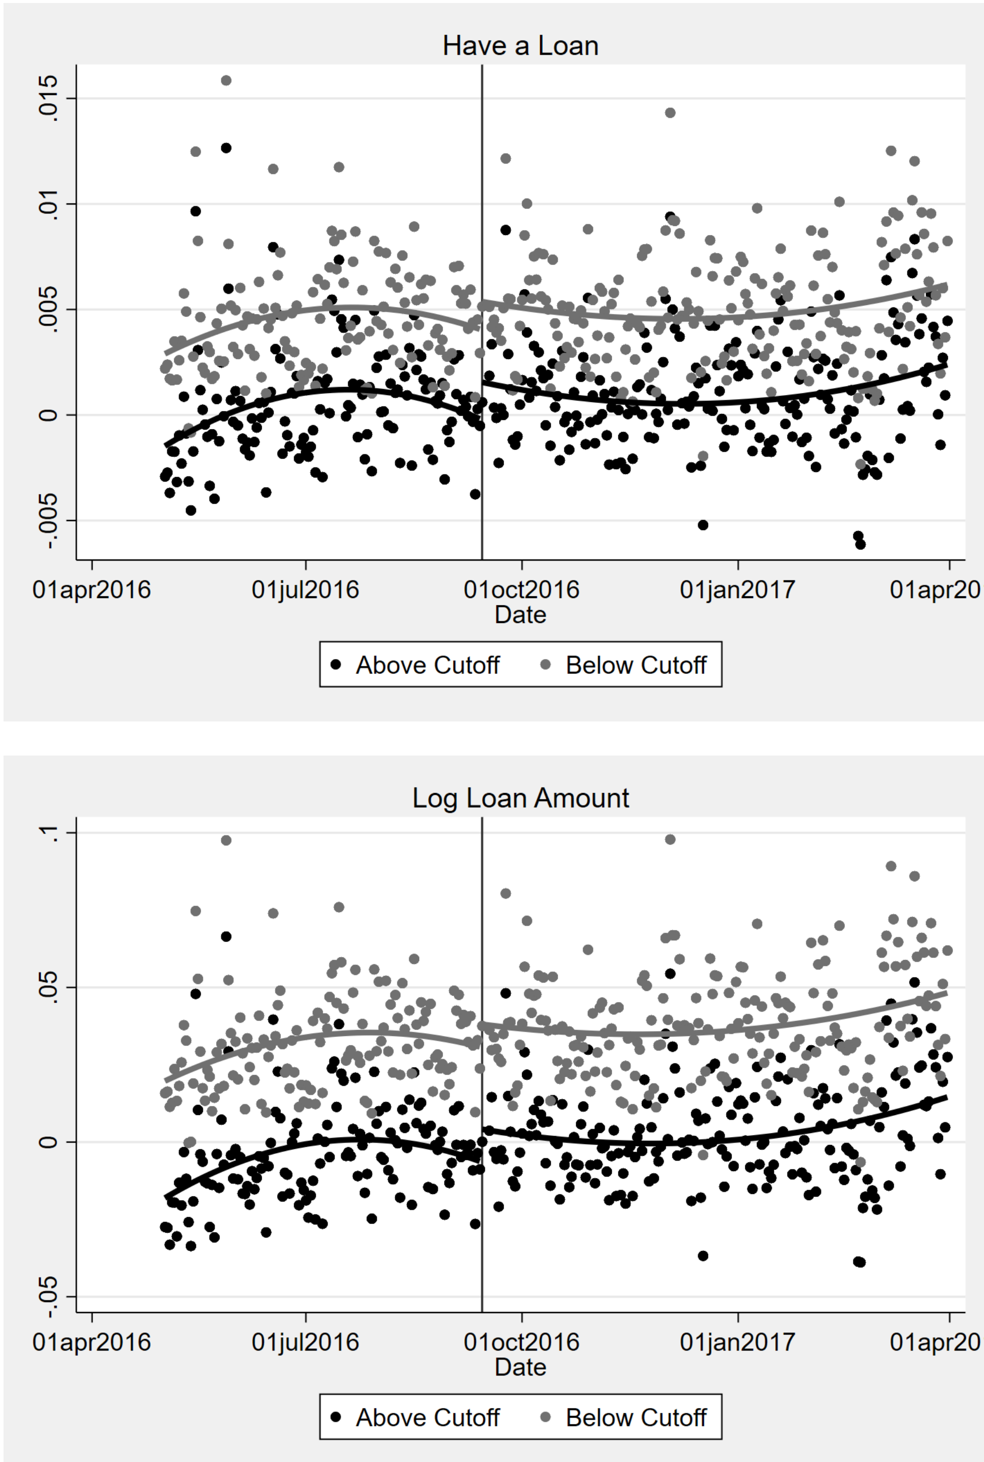

Table 1. Summary Statistics

	Mean	SD	N
Credit Score	48.196	63.507	899280
Post Interest Rate Caps	.489	.5	899370
Have a Loan	.018	.134	899370
Loan Amount	20.22	233.781	899370
Ln(1+Loan Amount)	.118	.878	899370
Limit	835.996	1243.498	899370
Fast Loans Pre Caps	.501	.5	700200
Limit Change Pre Caps	5.29	9.341	899370
Growing Limit Pre Caps	.628	.483	899370
Savings	623.601	6022.752	460000
Ln(1+Savings)	1.556	2.251	457369
Credit Score	.301	5.694	460000

Note: The loans data is daily data for a sample of 9,993 individuals for the optimal bandwidth window (46 days).
Fast loans refers to individuals whose speed between loan cycles was below the median.
Limit change refers to the actual change in the credit limit. A growing limit is a dummy for an increased limit.
The savings data is daily data for a sample of 5,000 individuals for the optimal bandwidth window (46 days).

Table 2. Impacts of Interest Rate Caps on Loans, Loan Amounts and Credit Limits

	Have a Loan				Log Loan Amount				Credit Limit			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Post Interest Rate Caps	.00185***	.00283***	.00235***	.00249***	.0113***	.0168***	.0146***	.0161***	8.41***	-40.4***	-48***	-90.6***
	[.000475]	[.000652]	[.000504]	[.000623]	[.00304]	[.00393]	[.00315]	[.00396]	[2.77]	[5.51]	[4.7]	[8.44]
Post Caps*Score Cutoff		-.00125**				-.00704**				62.4***		
		[.000591]				[.00347]				[6.69]		
Post Caps*Credit Score/100			-.00105***	-.000948**			-.00683***	-.00685**			117***	147***
			[.000367]	[.000463]			[.00247]	[.00316]			[9.23]	[12.6]
F-Statistic, Overall Effect		0.002	-0.048	-0.062		0.010	-0.315	-0.450		21.952	5595.971	9893.052
P-Value		0.002	0.006	0.047		0.002	0.008	0.035		0.000	0.000	0.000
Control Mean	0.018	0.018	0.018	0.019	0.119	0.119	0.119	0.124	766.937	766.937	767.014	875.301
Positive Score Sample				X				X				X
Observations	899370	899370	899280	704070	899370	899370	899280	704070	899370	899370	899280	704070

Note: * p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered at the individual level.

All regressions control for the running variable (date) and its interaction with the cutoff.

All specifications restricted to the CCT optimal bandwidth (46 days).

Table 3. Impacts of Interest Rate Caps, by Credit Score Quartile

	Have a Loan		Log Loan Amount		Credit Limit	
	(1)	(2)	(3)	(4)	(5)	(6)
Post Interest Rate Caps	0.0027***	0.0025*	0.0162***	0.0148**	-42.2481***	-65.0770***
	[0.0006]	[0.0013]	[0.0037]	[0.0073]	[5.4389]	[16.3682]
Post Caps*Score Quartile 2	-0.0008	-0.0003	-0.0036	-0.0004	-10.2917	-0.1890
	[0.0007]	[0.0013]	[0.0040]	[0.0073]	[7.1567]	[17.0783]
Post Caps*Score Quartile 3	-0.0009	-0.0004	-0.0046	-0.0014	55.8615***	65.9642***
	[0.0007]	[0.0013]	[0.0041]	[0.0073]	[9.3739]	[18.1196]
Post Caps*Score Quartile 4	-0.0019***	-0.0013	-0.0118***	-0.0086	159.1306***	169.2333***
	[0.0007]	[0.0013]	[0.0042]	[0.0074]	[12.5713]	[19.9622]
Credit Score Quartile 2	0.0013*	0.0032***	0.0079*	0.0206***	-28.3097**	79.5976***
	[0.0007]	[0.0012]	[0.0041]	[0.0071]	[14.0967]	[23.3361]
Credit Score Quartile 3	0.0013**	0.0032***	0.0221***	0.0349***	428.7154***	536.6227***
	[0.0007]	[0.0012]	[0.0043]	[0.0072]	[21.1455]	[28.1602]
Credit Score Quartile 4	0.0041***	0.0060***	0.0569***	0.0696***	1254.9627***	1362.8700***
	[0.0007]	[0.0012]	[0.0045]	[0.0074]	[31.6670]	[36.7244]
F-Statistic, Quartile 4 Effect	0.001	0.001	0.004	0.006	116.883	104.156
P-Value	0.162	0.082	0.293	0.170	0.000	0.000
Control Mean	0.018	0.019	0.119	0.124	766.937	875.189
Positive Score Sample		X		X		X
Observations	899370	704160	899370	704160	899370	704160

Note: * p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered at the individual level.
All regressions control for the running variable (date) and its interaction with the cutoff.
All specifications restricted to the CCT optimal bandwidth (46 days).

Table 4. Impacts of Interest Rate Caps on Savings

	Any Savings				Log Savings			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post Interest Rate Caps	0.0155***	0.0216***	0.0159***	0.0196**	0.0358*	0.0797***	0.0379**	-0.0044
	[0.0028]	[0.0039]	[0.0028]	[0.0096]	[0.0191]	[0.0263]	[0.0191]	[0.0527]
Post Caps*Score Cutoff		-0.0123**				-0.0889**		
		[0.0060]				[0.0357]		
Post Caps*Credit Score/100			-0.1068**	-0.0900			-0.7278**	0.4223
			[0.0529]	[0.1658]			[0.3145]	[0.8702]
Control Mean	0.809	0.809	0.809	0.761	1.577	1.577	1.577	1.614
Positive Score Sample				X				X
Observations	460000	460000	460000	227516	457369	457369	457369	226028

Note: * p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered at the individual level.

All regressions control for the running variable (date) and its interaction with the cutoff.

All specifications restricted to the CCT optimal bandwidth (46 days).

Table 5. Structural Estimation of Main Model Parameters

Parameter	**Estimate**	**Standard Error**
Sensitivity of Demand, β	0.001	0.0004
Latent Value, μ	0.248	0.1087
Dispersion Parameter, σ	0.011	0.2388

Table 6. Borrower Welfare and Average Market Shares in the Counterfactual

Scenario	Overall Average	Low Credit Score Average.	High Credit Score Average
Welfare			
Pre (baseline)	1.00	0.86	1.15
Post (actual)	0.97	0.79	1.16
Post (uniform)	0.62	0.00	1.30
Market Shares			
Pre (baseline), M-Shwari		0.50	0.65
Pre (baseline), Competitor		0.07	0.03
Post (actual), M-Shwari		0.52	0.60
Post (actual), Competitor		0.00	0.07
Post (uniform), M-Shwari		0.00	0.69
Post (uniform), Competitor		0.00	0.03

Table 7. Sensitivity of Welfare to σ

Scenario	$\sigma = 0.00$	$\sigma = 0.05$	$\sigma = 0.10$	$\sigma = 0.20$	$\sigma = 0.30$	$\sigma = 0.50$
Overall						
Pre (baseline)	1.00	1.01	1.02	1.07	1.14	1.33
Post (actual)	0.97	0.97	0.98	1.03	1.10	1.29
Post (uniform)	0.62	0.62	0.63	0.67	0.72	0.85
Low Credit Score Borrowers						
Pre (baseline)	0.87	0.87	0.88	0.90	0.94	1.05
Post (actual)	0.79	0.80	0.80	0.83	0.87	0.98
Post (uniform)	0.00	0.00	0.00	0.00	0.00	0.00
High Credit Score Borrowers						
Pre (baseline)	1.15	1.16	1.18	1.25	1.36	1.63
Post (actual)	1.16	1.16	1.18	1.26	1.36	1.63
Post (uniform)	1.30	1.31	1.33	1.40	1.50	1.77

Notes: Each column reports average borrower welfare (log-inclusive value) for a given value of the random coefficient dispersion σ, holding β and μ fixed at their estimated values. The welfare ranking across scenarios is preserved for all values of σ.

ONLINE APPENDIX

Appendix: Proofs

Proof of 1. By Equation 4, firm j serves market r if and only if $r \geq \underline{r}_j$, where $p(\underline{r}_j, 0) = c/i_j$. Since $i_1 > i_2$, we have $c/i_1 < c/i_2$, and since p is strictly increasing in r, it follows that $\underline{r}_1 < \underline{r}_2$. Therefore, for any $r \in [\underline{r}_1, \underline{r}_2)$, firm 1 finds it strictly profitable to lend while firm 2 does not. These are precisely the monopoly markets, and they are served exclusively by firm 1. □

Proof of 2. Throughout, $F_j(r, m_1, m_2)$ denotes the left-hand side of Equation 3 for firm j, and we assume the equilibrium is an interior, stable Nash equilibrium so that the second-order conditions hold ($\partial F_j/\partial m_j < 0$).

Monopoly markets. When only firm j is active, $s_j = 1$ and $\partial s_j/\partial m_j = 0$, so the FOC reduces to

$$G(r,m) \equiv i\,p(r,m) + i\,m\,\frac{\partial p(r,m)}{\partial m} - c = 0.$$

By the Implicit Function Theorem,

$$\frac{\partial \bar{m}}{\partial r} = -\frac{\partial G/\partial r}{\partial G/\partial m} = -\frac{i\,p_r + i\,m\,p_{rm}}{2i\,p_m + i\,m\,p_{mm}}.$$

The numerator is positive since $p_r > 0$ and $p_{rm} > 0$ (supermodularity). The denominator is negative since $p_m < 0$ and $p_{mm} \leq 0$ (p is decreasing and concave in m). Hence $\partial \bar{m}/\partial r > 0$.

Competitive markets. Applying the IFT to the system of FOCs for both firms simultaneously:

$$\frac{\partial \bar{m}_1(r)}{\partial r} = -\frac{\frac{\partial F_1}{\partial r} - \frac{\frac{\partial F_1}{\partial m_2}}{\frac{\partial F_2}{\partial m_2}}\frac{\partial F_2}{\partial r}}{\frac{\partial F_1}{\partial m_1} - \frac{\frac{\partial F_1}{\partial m_2}}{\frac{\partial F_2}{\partial m_2}}\frac{\partial F_2}{\partial m_1}}, \tag{16}$$

where F_j is the left-hand side of Equation 3 for firm j. We sign each component in turn.

$\partial F_j/\partial m_j < 0$. This is the second-order condition at the profit maximum.

$\partial F_j/\partial r > 0$. In the logit specification, shares s_j depend on credit limits and preference parameters but not directly on r. Differentiating F_j with respect to r thus involves only the repayment-probability terms:

$$\frac{\partial F_j}{\partial r} = \left[i\,p_r + i\,m_j\,p_{rm}\right]s_j + \left[i\,p_r\right]m_j\,\frac{\partial s_j}{\partial m_j}.$$

Every component is positive: $p_r > 0$, $p_{rm} > 0$ (supermodularity), $s_j > 0$, and $\frac{\partial s_j}{\partial m_j} > 0$.

$\partial F_1/\partial m_2 > 0$. Since only the share terms depend on m_2,

$$\frac{\partial F_1}{\partial m_2} = \left[i_1\,p + i_1\,m_1\,p' - c\right]\frac{\partial s_1}{\partial m_2} + \left[i_1\,p - c\right]m_1\,\frac{\partial^2 s_1}{\partial m_1\,\partial m_2},$$

where $p = p(r, m_1)$ and $p' = \partial p(r, m_1)/\partial m_1$. Substituting $[i_1\,p + i_1\,m_1\,p' - c] = -[i_1\,p - c]\,m_1\,\frac{\partial s_1/\partial m_1}{s_1}$ from the FOC gives

$$\frac{\partial F_1}{\partial m_2} = [i_1 p - c]\, m_1 \left[-\frac{1}{s_1} \frac{\partial s_1}{\partial m_1} \frac{\partial s_1}{\partial m_2} + \frac{\partial^2 s_1}{\partial m_1 \partial m_2} \right].$$

For the logit with inside-share normalization $s_1 + s_2 = 1$, we have $\frac{\partial s_1}{\partial m_1} = \beta(v - i_1) s_1 s_2$, $\frac{\partial s_1}{\partial m_2} = -\beta(v - i_2) s_1 s_2$, and $\frac{\partial^2 s_1}{\partial m_1 \partial m_2} = -\beta^2 (v - i_1)(v - i_2) s_1 s_2 (1 - 2s_1)$. After substitution, the bracketed expression simplifies to

$$\beta^2 (v - i_1)(v - i_2) s_1 s_2 \big[s_2 - (1 - 2s_1) \big] = \beta^2 (v - i_1)(v - i_2) s_1^2 s_2,$$

where the last equality uses $s_1 + s_2 = 1$. To see that $[i_1 p - c] > 0$, note that $p' < 0$ implies $[i_1 p + i_1 m_1 p' - c] < [i_1 p - c]$; since the FOC requires these two brackets (weighted by positive terms) to sum to zero, the first must be negative and the second strictly positive. Since $m_1 > 0$, we conclude $\partial F_1 / \partial m_2 > 0$.

Assembly. Define $R \equiv \frac{\partial F_1/\partial m_2}{\partial F_2/\partial m_2}$. Since $\frac{\partial F_1}{\partial m_2} > 0$ and $\frac{\partial F_2}{\partial m_2} < 0$, we have $R < 0$. The numerator of (16) is

$$\underbrace{\frac{\partial F_1}{\partial r}}_{>0} - \underbrace{R}_{<0} \underbrace{\frac{\partial F_2}{\partial r}}_{>0} > 0.$$

The denominator of (16) equals $\det(J) \big/ \big(\partial F_2 / \partial m_2 \big)$, where J is the Jacobian of (F_1, F_2) with respect to (m_1, m_2). At a stable equilibrium, $\det(J) > 0$ (own effects dominate cross effects), and since $\partial F_2 / \partial m_2 < 0$, the denominator is negative. Therefore $\frac{\partial \bar{m}_1}{\partial r} = -\frac{(+)}{(-)} > 0$. $\square$

Proof of 3. By Equation 4, the cutoff $\underline{r}_j$ solves $i_j p(r, 0) - c = 0$. Since p is strictly increasing in r, this equation has a unique solution for each i_j, provided one exists. After the policy, i_2 decreases to $i_2' < i_2$ while i_1 is unchanged. By assumption, i_2' is high enough that a solution $\underline{r}_2'$ to $i_2' p(r, 0) = c$ exists (otherwise firm 2 exits entirely). Since $i_2' p(\underline{r}_2, 0) < i_2 p(\underline{r}_2, 0) = c$, we must have $\underline{r}_2' > \underline{r}_2$ to restore equality (using $p_r > 0$). For firm 1, $i_1' = i_1$ implies $\underline{r}_1' = \underline{r}_1$. $\square$

Proof of 4. The first part follows directly from 3: since $r < \underline{r}_2'$, firm 2 no longer serves these markets after the policy, so $\bar{m}_2'(r) = 0$, while $\bar{m}_2(r) > 0$ since $r > \underline{r}_2$.

For the second part, define $h(m) \equiv i_1 p(r, m) + i_1 m \frac{\partial p(r, m)}{\partial m} - c$. Note that

$$h'(m) = 2 i_1 \frac{\partial p}{\partial m} + i_1 m \frac{\partial^2 p}{\partial m^2} < 0,$$

since p is decreasing ($\frac{\partial p}{\partial m} < 0$) and concave ($\frac{\partial^2 p}{\partial m^2} \le 0$) in m.

Pre-period (duopoly). Rearranging Equation 3 for firm 1:

$$h(\bar{m}_1) = -\big[i_1 p(r, \bar{m}_1) - c \big] \bar{m}_1 \frac{\partial s_1 / \partial m_1}{s_1}.$$

The right-hand side is strictly negative, since $i_1 p - c > 0$ (the firm earns positive margins), $\bar{m}_1 > 0$, and $\frac{\partial s_1}{\partial m_1} > 0$. Hence $h(\bar{m}_1) < 0$.

Post-period (monopoly). With firm 2 absent, $s_1 = 1$ and $\frac{\partial s_1}{\partial m_1} = 0$, so the FOC reduces to $h(\bar{m}_1') = 0$.

Because $r > \underline{r}_1$, we have $h(0) = i_1 p(r,0) - c > 0$. Since h is strictly decreasing, this guarantees a unique interior solution $\bar{m}_1' > 0$.

Since h is strictly decreasing and $h(\bar{m}_1) < 0 = h(\bar{m}_1')$, it follows that $\bar{m}_1 > \bar{m}_1'$. □

Proof of 5. In period t, the borrower's FOC is

$$\beta \frac{\partial \bar{m}_{1,t}}{\partial r}(r + a_{1,t}^*)(v - i_1) = c_{1,t}'(a_{1,t}^*).$$

Evaluate the post-period net marginal benefit at the pre-period optimum $a_{1,\text{pre}}^*$:

$$\begin{aligned}
&\beta \frac{\partial \bar{m}_{1,\text{post}}}{\partial r}(r + a_{1,\text{pre}}^*)(v - i_1) - c_{1,\text{post}}'(a_{1,\text{pre}}^*) \\
&= \underbrace{\beta \left[\frac{\partial \bar{m}_{1,\text{post}}}{\partial r} - \frac{\partial \bar{m}_{1,\text{pre}}}{\partial r} \right](v - i_1)}_{\text{change in marginal benefit}} + \underbrace{c_{1,\text{pre}}'(a_{1,\text{pre}}^*) - c_{1,\text{post}}'(a_{1,\text{pre}}^*)}_{\Delta c'(a_{1,\text{pre}}^*) > 0},
\end{aligned}$$

where the second equality uses the pre-period FOC. If the marginal benefit does not decrease, the first term is nonnegative and the expression is strictly positive. If the marginal benefit decreases, the stated condition ensures that $\Delta c'(a_{1,\text{pre}}^*)$ strictly exceeds the decrease, so the expression is again strictly positive. Since the post-period objective is strictly concave by assumption, a strictly positive net marginal benefit at $a_{1,\text{pre}}^*$ implies $a_{1,\text{post}}^* > a_{1,\text{pre}}^*$. □

Appendix: Identification of Structural Parameters

We formalize the identification argument for $\theta = (\beta, \mu, \sigma)$ from the moment conditions described in Section 6.3. Let $\mathscr{R}^+ = \{r : r > \hat{r}\}$ denote the set of competitive markets where firm 2 remains active in both periods.

Assumption 2. *(Time-invariance.) For each firm $j \in \{1,2\}$ and market r, the demand unobservable $\xi_{j,r}$ is the same in the pre- and post-policy periods.*

Assumption 3. *(Exclusion and relevance.) The instruments satisfy $\mathbb{E}[z_r \cdot \xi_{j,r}] = 0$ and $\mathbb{E}[z_r \cdot \bar{m}_{j,r}] \neq 0$ for at least one $j \in \{1,2\}$.*

Assumption 4. *(Support.) In $\mathscr{R}^+$, the credit limit $\bar{m}_{2,r}$ takes at least two distinct values.*

Proposition 6. *Under Assumptions 2–4, the parameters (β, μ, σ) are generically identified from the time-difference moment (11) and the instrumental variable moments (12).*

Proof. We proceed in three steps, using $\gamma \equiv \beta\sigma$ to denote the composite parameter that enters the share inversion.

Step 1 (Share inversion). For any $\gamma \geq 0$, the predicted share of firm j in market r and period t is

$$s_j(\delta_1, \delta_2; \gamma) = \int \frac{\exp(\delta_j + \gamma \bar{m}_j \nu)}{1 + \sum_{j'} \exp(\delta_{j'} + \gamma \bar{m}_{j'} \nu)} \, \phi(\nu) \, d\nu.$$

This is a smooth function that is strictly increasing in δ_j and strictly decreasing in δ_{-j}, with $s_j \to 0$ as $\delta_j \to -\infty$ and $s_j \to 1$ as $\delta_j \to \infty$. By standard contraction-mapping arguments, the mapping $(\delta_1, \delta_2) \mapsto (s_1, s_2)$ is a bijection for any fixed γ, so the observed shares uniquely determine mean utilities $\delta_{j,r,t}(\gamma) \equiv h_{j,r,t}(\gamma)$.

Step 2 (Identifying β and σ). By Assumption 2 and the definition $\delta_{j,r,t} = \beta \bar{m}_{j,r,t}(\mu - i_{j,t}) + \xi_{j,r}$, for any competitive market $r \in \mathscr{R}^+$ where $\bar{m}_{2,r}$ is constant across periods:

$$h_{2,r,1}(\gamma) - h_{2,r,0}(\gamma) = -\beta \, \bar{m}_{2,r} \, \Delta i_2.$$

This defines an implied sensitivity $\beta(r, \gamma) \equiv -[h_{2,r,1}(\gamma) - h_{2,r,0}(\gamma)]/(\bar{m}_{2,r} \, \Delta i_2)$. For the model to be internally consistent, $\beta(r, \gamma)$ must be the same across all $r \in \mathscr{R}^+$. By Assumption 4, there exist $r, r' \in \mathscr{R}^+$ with $\bar{m}_{2,r} \neq \bar{m}_{2,r'}$. Internal consistency therefore requires

$$\frac{h_{2,r,1}(\gamma) - h_{2,r,0}(\gamma)}{h_{2,r',1}(\gamma) - h_{2,r',0}(\gamma)} = \frac{\bar{m}_{2,r}}{\bar{m}_{2,r'}}.$$

The right-hand side is a known constant; the left-hand side is a nonlinear function of γ (through the dependence of the share inversion on γ). Generically—that is, for data not confined to a measure-zero set—this equation has a unique solution γ^*, yielding $\beta^* = \beta(r, \gamma^*)$ and $\sigma^* = \gamma^*/\beta^*$.

Step 3 (Identifying μ). Given (β^*, σ^*), the mean utilities are $\delta_{j,r} = h_{j,r}(\gamma^*)$, and $\xi_{j,r} = h_{j,r}(\gamma^*) - \beta^* \bar{m}_{j,r}(\mu - i_j)$. Assumption 3 requires:

$$\mathbb{E}\big[z_r \cdot \big(h_{j,r}(\gamma^*) - \beta^* \bar{m}_{j,r}(\mu - i_j)\big)\big] = 0,$$

which, given relevance ($\mathbb{E}[z_r \cdot \bar{m}_{j,r}] \neq 0$), uniquely determines

$$\mu = i_j + \frac{\mathbb{E}[z_r \cdot h_{j,r}(\gamma^*)]}{\beta^* \cdot \mathbb{E}[z_r \cdot \bar{m}_{j,r}]}.$$

$\square$

Remark 1. *The identification of σ in Step 2 relies on requiring internal consistency of the time-difference across markets with different credit limits. With only two inside options per market, this source of variation is weaker than in standard random-coefficients demand models with many products. In our estimates, this is reflected in the relatively large standard error on $\hat{\sigma}$.*

Figure A1. Digital Credit Market in Kenya, 2017

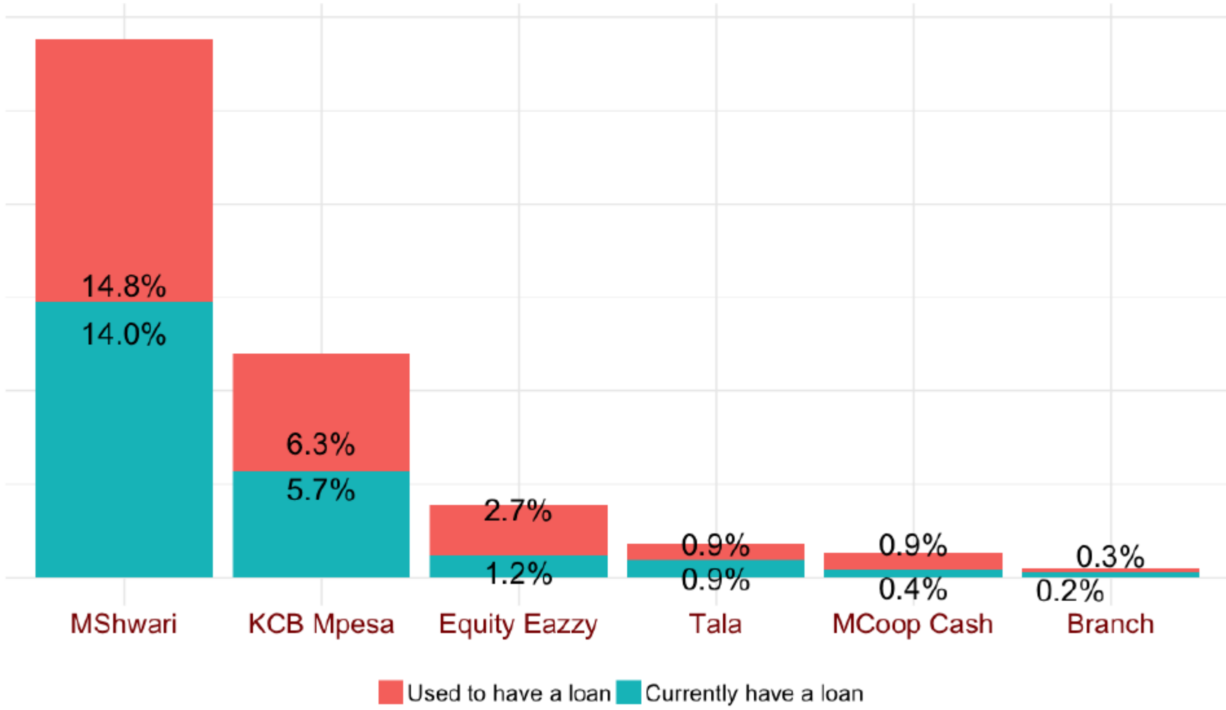

Note: This data shows the fraction of adults in Kenya who have borrowed from various digital credit platforms. Source: Totolo (2018).

Figure A2. Use of M-Shwari, 2017

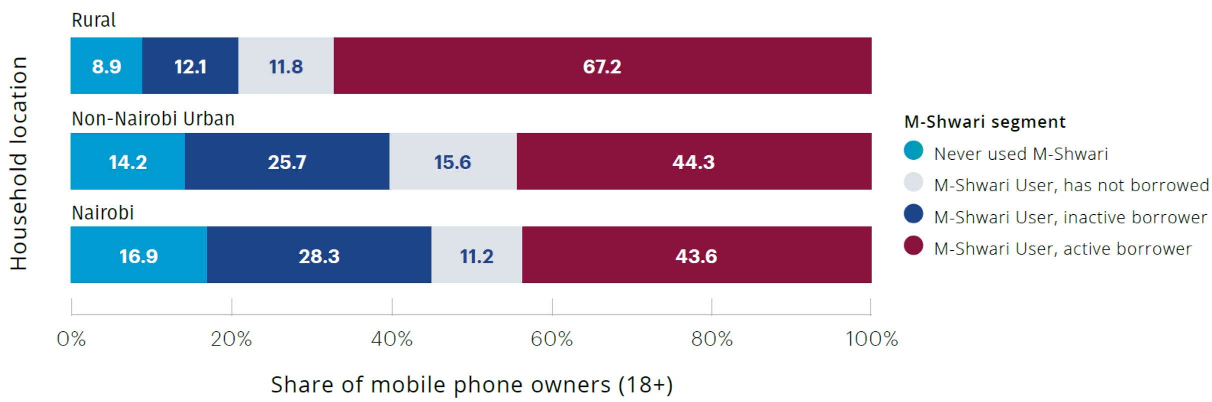

Note: This data shows how adults use M-Shwari in Kenya. Source: Suri and Gubbins (2018).

Table A1. Impacts of Interest Rate Caps on First Time Borrowers on M-Shwari

	First Loan			
	(1)	(2)	(3)	(4)
Post Interest Rate Caps	0.0007***	0.0004	0.0006***	0.0007***
	[0.0002]	[0.0003]	[0.0002]	[0.0002]
Post Caps*Score Cutoff		0.0005**		
		[0.0002]		
Post Caps*Credit Score/100			0.0002	0.0000
			[0.0001]	[0.0001]
F-Statistic, Overall Effect		0.001	0.010	0.004
P-Value		0.000	0.082	0.668
Control Mean	0.001	0.001	0.001	0.001
Positive Score Sample				X
Observations	899370	899370	899280	704070

Note: * p<0.1, ** p<0.05, *** p<0.01. Standard errors clustered at the individual level.
All regressions control for the running variable (date) and its interaction with the cutoff.
All specifications restricted to the CCT optimal bandwidth (46 days).