Cap On, Cap Off: Credit Rationing Hysteresis from Kenya's Interest Rate Ceiling

APEP Autonomous Research* @ai1scl

March 9, 2026

Abstract

Seventy-six countries regulate interest rates, yet credible evidence on their effects—and reversibility—remains scarce. I exploit Kenya's 2016 interest rate cap (CBR+4%) and its 2019 repeal using Central Bank supervisory data disaggregated by bank tier. In a difference-in-differences framework, small, SME-focused banks reduced their loan-to-asset ratio by 4.0 percentage points relative to large banks during the cap. After repeal, the gap widened to 6.5 percentage points—the differential contraction persisted and deepened rather than reversing. Government securities holdings remained elevated, and non-performing loan ratios stayed high. Cross-country comparisons with East African peers show consistent patterns. Randomization inference (within-year tier-label permutation) yields $p < 0.001$. The persistence is consistent with destroyed lending relationships and organizational lock-in. These findings suggest that even temporary interest rate ceilings may impose long-lasting costs on credit intermediation.

JEL Codes: G21, G28, O16, E43

Keywords: interest rate caps, credit rationing, hysteresis, financial regulation, Kenya, SME lending

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

In September 2016, Kenya capped interest rates on all bank loans at four percentage points above the Central Bank Rate. Three years later, Parliament abruptly repealed the law. This 38-month on-off experiment offers a rare window into a question with global stakes: when a government breaks a credit market with price controls, does it heal when the controls are removed?

At least 76 countries impose some form of interest rate ceiling on bank lending (Maimbo and Henriquez Gallegos, 2014; Ferrari et al., 2018). Stiglitz and Weiss (1981) showed that binding ceilings cause credit rationing because lenders cannot price risk and exit the riskiest segments. What remains poorly understood is whether these effects reverse when the cap is lifted. If credit rationing exhibits hysteresis—if temporary regulation causes persistent damage to lending capacity—then the costs of interest rate ceilings are far larger than standard welfare analysis suggests.

This paper studies that question using Kenya's Banking Amendment Act of 2016 and its repeal via the Finance Act of 2019. Most countries that impose rate ceilings never remove them, and when they do, confounding institutional changes make identification difficult. Kenya's cap was imposed abruptly, applied uniformly, and repealed cleanly, offering a valuable—though not perfectly clean—identification opportunity.

I exploit the differential bite of the cap across Kenya's tiered banking sector. The Central Bank of Kenya (CBK) classifies commercial banks into three peer groups: Tier 1 (six large, diversified banks), Tier 2 (roughly 15 medium banks), and Tier 3 (roughly 20 small, SME-focused banks). Before the cap, Tier 3 banks charged substantially higher lending rates to compensate for serving riskier clienteles. The cap compressed all banks to the same ceiling, but the constraint bound most tightly on small banks that had relied on risk-based pricing. I use this variation in a difference-in-differences framework, comparing outcomes for Tier 3 banks relative to Tier 1 banks across the pre-cap (2010–2016), cap-on (2017–2019), and post-repeal (2020–2023) periods.

The main finding is that the differential contraction does not reverse. Small banks reduced their loan-to-asset ratio by 4.0 percentage points relative to large banks during the cap (RI $p < 0.001$). After repeal, the gap *widened* to 6.5 percentage points. An event study shows flat pre-trends through $t - 2$, a sharp break at $t = 0$ (the cap year), and persistently negative coefficients through the end of the sample. Randomization inference, based on 1,000 within-year shuffles of tier treatment labels (sampling from $3^{14} \approx 4.8$ million possible assignment vectors), yields $p < 0.001$ for all four main outcomes.

Three channels are consistent with this persistence. First, banks rebalanced portfolios toward government securities. Tier 3 banks' government securities share rose from 26.7% pre-cap to 33.8% during the cap and continued climbing to 37.2% post-repeal. The risk-free yield on Kenyan Treasury bills made this substitution attractive, and once banks built the institutional capacity for sovereign debt trading, they had little reason to return to costly SME relationship lending. Second, lending relationships were likely disrupted. SME borrowers cut off during the cap period may have found alternative (often informal or digital) credit sources, and the bank-specific soft information accumulated through years of relationship lending would have dissipated—though I cannot observe this directly in tier-level data. Third, the non-performing loan ratio for Tier 3 banks jumped from 9.0% to 13.5% during the cap and remained at 13.4% post-repeal, leaving banks risk-averse even after regulatory constraints were lifted.

Descriptive cross-country evidence provides suggestive corroboration. Using World Development Indicators for Kenya, Uganda, Tanzania, and Rwanda, Kenya's lending rate fell 3.4 percentage points relative to East African peers during the cap ($p = 0.075$), while its NPL ratio rose 4.4 percentage points ($p = 0.109$). Both patterns persist after repeal. With only four countries and limited statistical power, these results are best interpreted as descriptive context consistent with—but not independent confirmation of—the within-Kenya findings.

The welfare consequences extend beyond the formal banking sector. Kenya's world-leading mobile money ecosystem—anchored by M-Pesa—enabled rapid growth in digital credit products during the cap period. M-Shwari and Tala users grew from approximately 200,000 to 2 million between 2016 and 2019, offering instant mobile loans at effective annual rates exceeding 138% (Bharadwaj et al., 2021; Suri and Jack, 2016). Borrowers rationed out of formal banking did not stop borrowing; they appear to have shifted to vastly more expensive channels. While I cannot establish a causal link between the cap and digital credit growth with the available data, the coincident timing suggests a substitution channel that conventional analysis—focused on the formal sector—would miss.

This paper contributes to several literatures. Most directly, it advances the empirical literature on interest rate ceilings, which has been surprisingly thin on credible causal evidence. Degryse et al. (2009) survey the theory but note the scarcity of clean natural experiments. The closest empirical predecessor is Igan and Pinheiro (2023), who study the Kenyan cap using aggregate data and find reduced credit growth, but cannot identify the hysteresis effect because they focus on the cap-on period alone. Maimbo and Henriquez Gallegos (2014) provide a cross-country survey of interest rate regulations, and Ferrari et al. (2018) examine the global landscape, but neither exploits a cap-and-repeal episode for causal identification. This paper provides the first evidence that credit rationing effects from interest rate ceilings may persist well beyond the policy's removal.

Second, the paper contributes to the broader literature on financial regulation and its unintended consequences (Agarwal et al., 2014; Buchak et al., 2018; D'Acunto and Rossi, 2022). The finding that regulatory damage may persist after the regulation is removed echoes work on regulatory uncertainty in other settings (Baker et al., 2016) and provides suggestive evidence from a symmetric on-off design.

Third, the paper speaks to the growing literature on relationship lending and small-business credit in developing countries (Beck et al., 2005; Banerjee et al., 2015; Petersen and Rajan, 1994). The persistence pattern—consistent with destroyed lending relationships that do not spontaneously regenerate—aligns with theoretical models of relationship lending (Boot and Thakor, 2000; Bolton et al., 2016) but has not previously been documented in a policy-induced setting.

Finally, the paper contributes to the fintech literature by documenting a plausible channel through which banking regulation may push borrowers toward alternative credit markets (Jack and Suri, 2014; Suri and Jack, 2016; Bharadwaj et al., 2021; Mbiti and Weil, 2016). The coincident growth in digital credit during the cap period is consistent with the cap acting as an inadvertent subsidy to digital lenders, whose unregulated rates eclipsed anything banks had charged.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting, including the structure of Kenya's banking sector, the cap legislation, and the digital credit ecosystem. Section 3 develops the conceptual framework linking adverse selection to portfolio rebalancing and hysteresis. Section 4 describes the data. Section 5 lays out the empirical strategy. Section 6 presents the main results and event study. Section 7 examines mechanisms. Section 8 provides robustness checks. Section 9 discusses implications. Section 10 concludes.

2. Institutional Background

2.1 Kenya's Banking Sector

Kenya operates one of Sub-Saharan Africa's most developed financial systems. As of 2016, the sector comprised 42 commercial banks regulated by the Central Bank of Kenya (CBK), serving a population of approximately 48 million. Total banking sector assets exceeded KES 3.7 trillion (roughly USD 37 billion), with a private credit-to-GDP ratio of about 32%—high by regional standards but below the levels typical of middle-income countries (Beck et al., 2005; Demirgüç-Kunt et al., 2012).

The CBK classifies commercial banks into three peer groups based on market share of net assets, a classification that captures meaningful structural differences in business models

(Central Bank of Kenya, 2017). *Tier 1* consists of the six largest banks (including Equity Bank, KCB Group, and Co-operative Bank), which collectively hold approximately 50% of sector assets. These banks serve large corporate clients and the salaried middle class, operate extensive branch networks, and maintain diversified revenue streams that include fee income, foreign exchange trading, and investment banking. Their pre-cap lending rates typically ranged from 14–18%, and they held substantial government securities portfolios as part of standard asset-liability management.

Tier 2 comprises roughly 15 mid-sized banks that blend retail and corporate lending. These banks were moderately affected by the cap because their pre-cap rates, while higher than Tier 1, were not dramatically above the ceiling. *Tier 3* consists of approximately 20 small banks that specialize in SME lending, trade finance for small enterprises, and community banking. These institutions charged the highest pre-cap lending rates—often 20–24%—because their clientele presented higher credit risk and because the fixed costs of small-ticket lending are substantial relative to loan size (Stiglitz and Weiss, 1981; Banerjee et al., 2015). The cap bound most tightly on these banks.

The tier classification is stable over time: banks rarely move between tiers, and the CBK's peer group assignment is based on structural characteristics rather than short-run performance. This stability is important for identification because it means the treatment intensity (exposure to the cap) was determined by pre-existing bank characteristics rather than by endogenous responses to the policy.

2.2 The Banking Amendment Act of 2016

On September 14, 2016, Kenya's Banking Amendment Act came into force, capping all bank lending rates at 4 percentage points above the Central Bank Rate (Were and Wambua, 2019). At the time of implementation, the CBR stood at 10%, implying a ceiling of 14% on all bank loans regardless of risk. The law also imposed a floor on deposit rates at 70% of the CBR, compressing the net interest margin from both directions.

The motivation was populist: Kenyan politicians had long complained that banks charged "usurious" rates while earning record profits (Mbugua and Rotich, 2018). The cap followed a failed voluntary initiative by the Kenya Bankers Association to reduce lending rates, and it passed with overwhelming parliamentary support despite vigorous opposition from the CBK, the International Monetary Fund, and the banking industry (International Monetary Fund, 2018).

Several features of the cap matter for identification. First, it was *uniform*: all banks, all loan products, and all borrower types faced the same ceiling. There were no exemptions for microfinance, agricultural lending, or small-ticket loans. Second, compliance was effectively perfect: the CBK monitored adherence through regular reporting, and violations would have risked banking licenses. Third, the cap was *unanticipated* in its timing—while the debate had simmered for years, the specific bill passed rapidly through Parliament in August–September 2016, giving banks minimal time to adjust in advance. Fourth, the cap was *simple*: it was defined relative to the CBR, a single publicly observable benchmark, leaving no room for creative compliance through fee restructuring or product redesign (Igan and Pinheiro, 2023).

2.3 The Finance Act of 2019 and Repeal

On November 7, 2019, President Uhuru Kenyatta signed the Finance Act 2019, which repealed the interest rate cap effective immediately. The repeal followed sustained pressure from the IMF, which had made removal of the cap a condition for its precautionary credit facility (International Monetary Fund, 2018), and from the CBK, whose governor publicly argued that the cap was choking private sector credit growth (Central Bank of Kenya, 2019).

The repeal was similarly clean from an identification perspective. It applied immediately to all new loans (though existing loan contracts at the capped rate were honored through maturity). There were no phase-out provisions, no replacement regulations, and no simultaneous banking sector reforms that might confound estimation. The timing was largely driven by the external pressure from the IMF credit facility rather than by domestic macroeconomic conditions, which aids identification.

The 38-month duration of the cap (September 2016 through November 2019) provides sufficient variation for both the cap-on and the post-repeal period, with data through 2023 allowing four full years of post-repeal observation. Figure 1 presents the policy timeline.

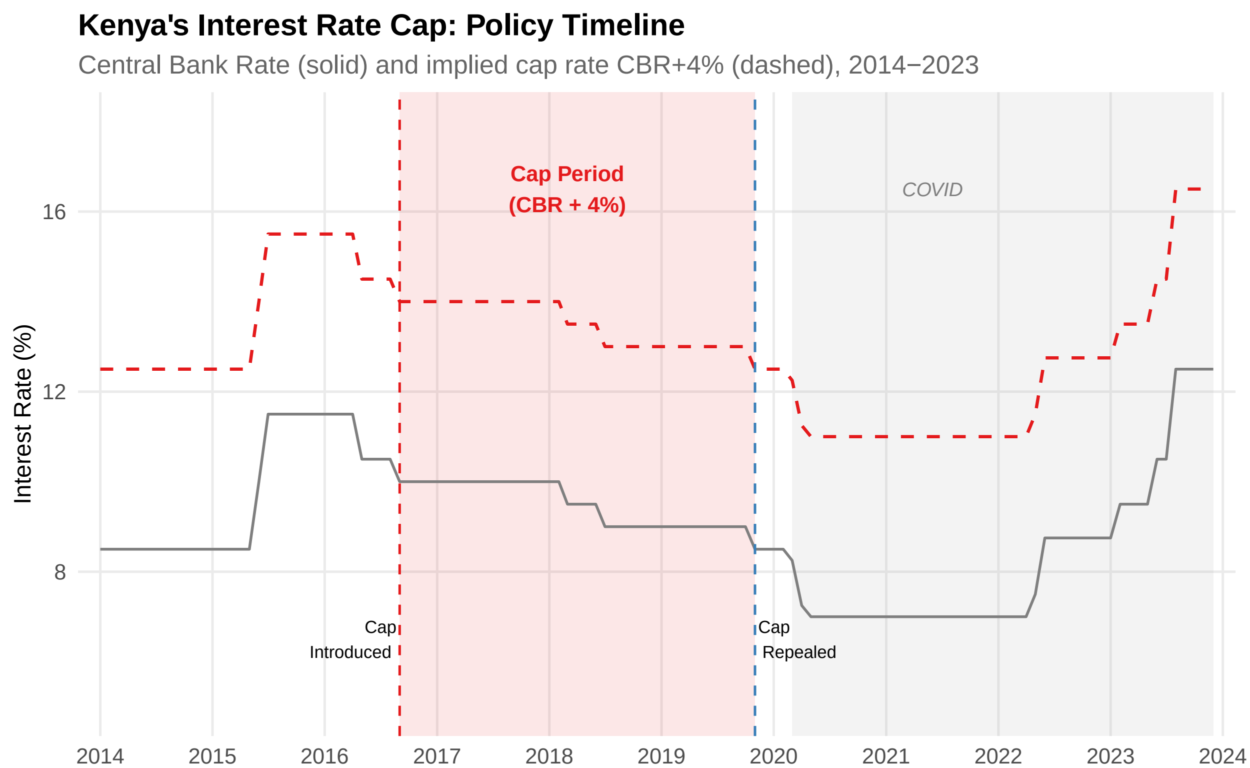

Figure 1: Policy Timeline: Interest Rate Cap and Repeal

Notes: The Banking Amendment Act took effect September 14, 2016, capping lending rates at CBR+4%. The Finance Act 2019 repealed the cap on November 7, 2019. Vertical dashed lines indicate policy transitions.

2.4 The M-Pesa Ecosystem and Digital Credit

Kenya's financial landscape is unusual among developing countries in the depth of its mobile money infrastructure. Safaricom's M-Pesa, launched in 2007, had reached near-universal adoption by the time the cap was imposed, with over 25 million active users and annual transaction volumes exceeding 50% of GDP (Jack and Suri, 2014; Suri and Jack, 2016). This infrastructure enabled the rapid deployment of digital credit products that became a critical substitute when bank lending contracted.

M-Shwari, a joint product of the Commercial Bank of Africa (now NCBA) and Safaricom, launched in 2012 and offered instant savings and loan products through the M-Pesa interface. By 2016, it had approximately 200,000 active borrowers. During the cap period, as banks rationed credit to SMEs and informal businesses, digital credit platforms including M-Shwari, KCB M-Pesa, Tala, and Branch experienced explosive growth. By 2019, digital credit users had surpassed 2 million (Bharadwaj et al., 2021; Totolo, 2018).

The pricing structure of these digital products starkly illustrates the welfare cost of the cap. M-Shwari charges a "facilitation fee" of 7.5% per 30-day loan period, translating to an effective APR of approximately 138%. Tala and Branch charge similar rates. These rates are roughly ten times the regulated bank rate and substantially higher than the pre-cap unregulated bank rates they replaced. The cap, designed to make credit cheaper, channeled borrowers into the most expensive credit market in Kenya's history (Totolo, 2018; Mbiti and Weil, 2016).

3. Conceptual Framework

The conceptual framework proceeds in three stages: how the cap causes credit rationing (drawing on Stiglitz and Weiss 1981), how rationed banks rebalance portfolios, and why the resulting equilibrium exhibits hysteresis.

3.1 Adverse Selection and Credit Rationing Under a Ceiling

Consider a bank that serves a pool of borrowers with heterogeneous risk types $\theta \sim F(\theta)$, where higher θ indicates greater default probability. In the absence of regulation, the bank sets its lending rate r^* to maximize expected profit:

$$\pi(r) = (1 - p(r)) \cdot r \cdot L - c(L) \tag{1}$$

where $p(r)$ is the average default probability (endogenous to r through adverse selection), L is the loan volume, and $c(\cdot)$ captures funding and operational costs. Stiglitz and Weiss (1981) showed that $p(r)$ is increasing in r because higher rates drive safe borrowers out of the market (adverse selection) and encourage risk-taking by remaining borrowers (moral hazard). The profit-maximizing rate r^* trades off higher revenue per loan against deteriorating portfolio quality.

Now impose a ceiling $\bar{r} < r^*$. The bank cannot price risk for its riskiest borrowers, so it must either lend to them at a loss (since $\bar{r}$ does not compensate for their default probability) or refuse them credit. Rational banks choose the latter: they ration credit by tightening non-price terms (collateral requirements, loan size limits) or by exiting risky market segments entirely.

Prediction 1: Banks with higher pre-cap exposure (higher $r^* - \bar{r}$) reduce lending more. This maps directly to the Tier 3 vs. Tier 1 comparison: Tier 3 banks charged the highest pre-cap rates and therefore faced the largest gap between their optimal and regulated rates.

3.2 Portfolio Rebalancing Toward Government Securities

Banks that reduce private lending do not simply shrink their balance sheets. Kenya's Treasury bill and bond market offers risk-free yields of 8–12%, providing an attractive alternative for assets freed from the loan portfolio. Define the bank's asset allocation problem:

$$\max_{L,G} \quad (1 - p(\bar{r})) \cdot \bar{r} \cdot L + r_G \cdot G - c(L + G) \tag{2}$$

subject to $L + G \leq A$ (total assets), where G is government securities and r_G is the sovereign yield. When the cap reduces the risk-adjusted return on private lending, the optimal portfolio tilts toward G. The strength of this substitution depends on r_G relative to the constrained expected return on lending—and in Kenya, sovereign yields were sufficiently high to make the substitution attractive even before the cap.

Prediction 2: Banks with higher cap exposure increase government securities holdings. The magnitude of substitution should be largest for Tier 3 banks, which face the biggest decline in risk-adjusted lending returns.

3.3 Hysteresis Through Relationship Destruction

The key theoretical contribution is explaining why the portfolio rebalancing does not reverse when the cap is removed. Three channels generate hysteresis:

Channel 1: Relationship capital destruction. Relationship lending produces borrower-specific soft information through repeated interactions (Petersen and Rajan, 1994; Boot and Thakor, 2000; Bolton et al., 2016). When a bank exits an SME lending segment, the accumulated relationship capital depreciates. Loan officers with SME expertise leave or are reassigned. Credit scoring models for small borrowers become stale. After three years of attrition, rebuilding this capacity requires costly investment with uncertain returns. Since the bank has already built capacity for government securities trading, the opportunity cost of returning to SME lending has risen.

Channel 2: Borrower displacement. SME borrowers cut off from formal bank credit during the cap period found alternatives—digital credit, informal moneylenders, retained earnings, or simply reduced investment. Upon repeal, these borrowers do not automatically return to banks: some have established alternative credit relationships, some have suffered financial distress and no longer qualify, and some have downsized their businesses. The demand for formal SME credit does not snap back.

Channel 3: Organizational learning lock-in. Banks that rebalanced toward government securities invested in the institutional infrastructure for sovereign debt management—treasury desks, risk management systems, regulatory compliance for securities trading. This creates path dependence: the bank's comparative advantage has shifted, and internal incentives now favor the new portfolio composition (Hannan and Freeman, 1984).

Prediction 3: The cap effect does not reverse upon repeal. Formally, the symmetry test is $H_0 : \hat{\beta}_{\text{Repeal}} = 0$ (full reversal, outcomes return to baseline) versus $H_1 : \hat{\beta}_{\text{Repeal}} < 0$ (hysteresis, damage persists). Under strong hysteresis, the post-repeal coefficient may even be more negative than the cap-on coefficient ($\hat{\beta}_{\text{Repeal}} < \hat{\beta}_{\text{Cap}} < 0$), implying the damage deepens as relationship capital continues to depreciate.

4. Data

4.1 Central Bank of Kenya Supervisory Data

The primary data source is the Central Bank of Kenya's Annual Bank Supervision Reports, which publish balance sheet and income statement data disaggregated by the CBK's peer group (tier) classification (Central Bank of Kenya, 2017, 2019). I compile data for the period 2010–2023, covering seven pre-cap years (2010–2016), three cap-on years (2017–2019), and four post-repeal years (2020–2023).

The key variables are: (i) total assets, (ii) loans and advances to customers, (iii) government securities holdings, (iv) gross non-performing loans, and (v) total deposits. From these, I construct three outcome ratios: the loan-to-asset ratio (loans and advances divided by total assets), the government securities-to-asset ratio, and the NPL ratio (non-performing loans as a share of total loans). These ratios are the standard metrics for bank portfolio composition used in the banking regulation literature (Degryse et al., 2009).

The unit of observation is the tier-year. Because the CBK reports aggregate data at the tier level rather than individual bank data, the panel has three tiers $\times$ 14 years = 42 observations. This is a limitation—individual bank data would permit more granular analysis—but the tier-level aggregates are the most reliable publicly available data, and the tier classification captures the primary source of variation in cap exposure.

4.2 World Development Indicators

For the cross-country analysis, I draw on the World Bank's World Development Indicators (WDI), which provide annual data on domestic credit to the private sector as a share of GDP, average lending interest rates, and bank non-performing loan ratios for Kenya and three East African comparators: Uganda, Tanzania, and Rwanda (World Bank, 2024). These countries share Kenya's East African Community membership, similar macroeconomic environments, and regulatory traditions, but did not impose interest rate caps during the study period. The cross-country panel spans 2010–2023 with four countries, yielding approximately 45 country-year observations after accounting for missing data.

4.3 Digital Credit Data

Data on digital credit adoption come from industry reports and surveys compiled by the Financial Sector Deepening Trust (FSD Kenya), the Consultative Group to Assist the Poor (CGAP), and academic studies (Totolo, 2018; Bharadwaj et al., 2021). These sources provide estimates of digital credit user counts, loan volumes, and pricing structures for the major platforms (M-Shwari, KCB M-Pesa, Tala, Branch). The digital credit data are used descriptively to document the substitution channel and estimate its welfare cost; they do not enter the formal econometric analysis.

4.4 Summary Statistics

Table 1 presents summary statistics by bank tier and policy period. Several patterns are immediately apparent. Tier 3 banks are an order of magnitude smaller than Tier 1 banks in total assets (KES 161 billion vs. KES 1,404 billion pre-cap). Despite their small size, Tier 3 banks maintained comparable loan-to-asset ratios before the cap (0.552 vs. 0.573), reflecting their SME-focused business model. The divergence during and after the cap is stark: Tier 3's loan-to-asset ratio fell from 0.552 to 0.465 (cap-on) and then to 0.436 (post-repeal), while Tier 1's ratio declined modestly from 0.573 to 0.527 and stabilized at 0.524.

The government securities share tells the mirror story: Tier 3 banks increased their sovereign debt holdings from 26.7% to 33.8% (cap-on) and 37.2% (post-repeal), a 10.5 percentage point cumulative increase. NPL ratios rose sharply for Tier 3 (9.0% to 13.5%) and remained elevated post-repeal (13.4%), while Tier 1 NPLs rose modestly (5.0% to 8.2%) and stabilized.

The number of Tier 3 banks declined from approximately 22 (pre-cap) to 19 (cap-on) and 16 (post-repeal), reflecting consolidation as several small banks were placed under CBK receivership or merged. This compositional change, if anything, biases against finding hysteresis: the surviving Tier 3 banks are the strongest of their peer group, and yet their lending ratios remained depressed.

Table 1: Summary Statistics by Bank Tier and Policy Period

Tier	Period	Years	Assets (KES bn)	Loans (KES bn)	Loan/ Asset	Govt Sec/ Asset	NPL (%)
Tier 1 (Large, $n = 6$)							
	Pre-Cap (2010–2016)	7	1,404	807	0.573	0.184	5.0
	Cap On (2017–2019)	3	2,277	1,200	0.527	0.221	8.2
	Post-Repeal (2020–2023)	4	3,038	1,595	0.524	0.220	8.2
Tier 2 (Medium, $n \approx 15$)							
	Pre-Cap (2010–2016)	7	689	396	0.572	0.213	5.7
	Cap On (2017–2019)	3	1,040	547	0.526	0.272	10.0
	Post-Repeal (2020–2023)	4	1,308	680	0.520	0.272	10.2
Tier 3 (Small, $n \approx 21$)							
	Pre-Cap (2010–2016)	7	161	89	0.552	0.267	9.0
	Cap On (2017–2019)	3	204	95	0.465	0.338	13.5
	Post-Repeal (2020–2023)	4	218	95	0.436	0.372	13.4

Notes: Source: Central Bank of Kenya Annual Reports, 2010–2023. Tier classification follows CBK peer grouping. Loan/Asset is the ratio of loans and advances to total assets. Govt Sec/Asset is government securities holdings divided by total assets. NPL ratio is non-performing loans as a percentage of total loans.

5. Empirical Strategy

5.1 Identification

The identification strategy exploits differential exposure to the interest rate cap across bank tiers. The cap was uniform—all banks faced the same CBR+4% ceiling—but it bound differentially because banks entered the cap period with different lending rate structures. Tier 3 banks, which charged the highest pre-cap rates to compensate for serving riskier SME clienteles, experienced the largest reduction in feasible lending margins. Tier 1 banks, whose pre-cap rates were closer to the ceiling, were less constrained.

This differential exposure generates a natural treatment intensity that I exploit in a difference-in-differences framework. The identifying assumption is that Tier 3 and Tier 1 banks would have followed parallel trends in portfolio composition absent the cap. I provide several pieces of evidence supporting this assumption. First, the event study (Figure 2) shows that the Tier 3-vs-Tier 1 gap in loan-to-asset ratios is flat and statistically indistinguishable from zero for $t-6$ through $t-2$ (2010–2014), with no systematic pre-trend. Second, both tiers operated in the same macroeconomic environment, were subject to the same CBK regulations (aside from the cap), and faced the same external shocks. Third, a placebo test using the Tier 2 vs. Tier 1 comparison—where the cap's differential bite is minimal—finds no significant effects, as expected under the parallel trends assumption.

5.2 Estimation

The main specification is:

$$Y_{it} = \alpha_i + \delta_t + \beta_1(\text{Tier3}_i \times \text{CapOn}_t) + \beta_2(\text{Tier3}_i \times \text{PostRepeal}_t) + \varepsilon_{it} \tag{3}$$

where Y_{it} is the outcome (loan-to-asset ratio, government securities ratio, or NPL ratio) for tier i in year t; α_i are tier fixed effects; δ_t are year fixed effects; Tier3_i is an indicator for Tier 3 banks; $\text{CapOn}_t = \mathbb{I}[2017 \leq t \leq 2019]$; and $\text{PostRepeal}_t = \mathbb{I}[t \geq 2020]$. Standard errors are clustered at the tier level.

The coefficient β_1 captures the differential effect of the cap on Tier 3 relative to Tier 1 banks during the cap-on period. The coefficient β_2 captures the differential effect in the post-repeal period—both measured relative to the pre-cap baseline (the omitted category, 2010–2016). Under full reversal, $\beta_2 = 0$ (Tier 3 returns to baseline). Under full hysteresis, $\beta_2 \approx \beta_1$ (the damage persists). The data will reveal that $\beta_2 < \beta_1 < 0$—the damage deepens.

I also estimate an event-study specification:

$$Y_{it} = \alpha_i + \delta_t + \sum_{k \neq -1} \gamma_k(\text{Tier3}_i \times \mathbb{I}[t - t^* = k]) + \varepsilon_{it} \tag{4}$$

where $t^* = 2016$ (the partial cap year—the cap took effect in September 2016), k indexes years relative to the cap, and $k = -1$ (2015) is the omitted reference period. **Treatment timing note:** The cap was legislated in September 2016 and repealed in November 2019. With annual data, I classify 2016 as a transition year: it enters the pre-cap category in the main DiD (where $\text{CapOn}_t = \mathbb{I}[2017 \leq t \leq 2019]$ captures only full cap years) but serves as the $k = 0$ event-study year to show the immediate partial-year effect. Similarly, 2019 is coded as the last full cap year in the main DiD and as $k = 3$ in the event study. This conservative coding ensures the DiD coefficients reflect full-year exposure rather than partial-year effects. The coefficients γ_k trace out the dynamic treatment effect, allowing visual inspection of pre-trends and the post-repeal trajectory.

For the cross-country analysis, I estimate:

$$Y_{ct} = \alpha_c + \delta_t + \phi_1(\text{Kenya}_c \times \text{CapOn}_t) + \phi_2(\text{Kenya}_c \times \text{PostRepeal}_t) + \varepsilon_{ct} \tag{5}$$

where c indexes countries (Kenya, Uganda, Tanzania, Rwanda) and the outcomes are credit-to-GDP, lending rates, and NPL ratios from the WDI. Standard errors are clustered at the country level.

5.3 Symmetry Test

The symmetry test formalizes the hysteresis hypothesis. I define the reversal ratio as:

$$\text{Reversal} = -\frac{\hat{\beta}_2 - \hat{\beta}_1}{\hat{\beta}_1} \times 100\% \tag{6}$$

Under full reversal, $\hat{\beta}_2 = 0$ and Reversal $= 100\%$. Under full hysteresis (persistent damage), $\hat{\beta}_2 = \hat{\beta}_1$ and Reversal $= 0\%$. If the damage deepens post-repeal, Reversal $< 0\%$. This metric provides a transparent, unit-free summary of the policy's reversibility.

5.4 Threats to Validity

Several concerns warrant discussion. First, *compositional changes*: the number of Tier 3 banks fell from 22 to 16 over the sample period as weak banks were resolved. If exiting banks had higher loan-to-asset ratios, their departure would mechanically reduce the tier average. However, this biases against finding hysteresis because the surviving banks are the strongest in the peer group. Second, *COVID-19*: the pandemic began in March 2020, shortly after the cap's repeal, potentially confounding the post-repeal period. I address this by estimating a pre-COVID specification using 2010–2019 data, which isolates the cap effect from pandemic disruption. Third, *simultaneous policies*: no other major banking regulation was introduced alongside the cap or its repeal, though the CBK did adjust the CBR several times. Since the cap was defined relative to the CBR, CBR changes affected the cap level but not its existence—and my analysis uses portfolio composition ratios rather than interest rate levels. Fourth, *small sample*: the tier-year panel has only 42 observations. This is addressed through randomization inference, which does not rely on asymptotic approximations and instead constructs the exact null distribution.

6. Results

6.1 Main Difference-in-Differences Results

Table 2 reports estimates from the main specification (Equation (3)). Column (1) shows the loan-to-asset ratio. The Tier 3 × Cap On coefficient is −0.040 (RI $p < 0.001$), indicating that Tier 3 banks reduced their loan-to-asset ratio by 4.0 percentage points relative to Tier 1 banks during the cap period. Because the panel has only three tier-level clusters, cluster-robust standard errors are reported for scale but are known to be downward biased with few clusters (Cameron et al., 2008); all inferential claims for the main specification rest on randomization inference (Table 5). The Tier 3 × Post-Repeal coefficient is −0.065, indicating that Tier 3's loan-to-asset ratio remained 6.5 percentage points below Tier 1's relative to the pre-cap baseline—the gap not only persisted but widened beyond the 4.0 pp cap-on effect.

Column (2) shows government securities as a share of total assets. Tier 3 banks increased their sovereign debt holdings by 2.3 percentage points during the cap, and the substitution deepened to 5.8 percentage points post-repeal. This result is precisely the mirror image of the lending contraction: assets freed from the loan portfolio flowed into government securities.

Column (3) reports the NPL ratio. The point estimates are positive (0.83 percentage points during the cap, 0.59 pp post-repeal), consistent with deteriorating portfolio quality. The cluster-robust SEs (0.62 and 0.72) are large relative to the coefficients, but as noted above, these SEs are unreliable with 3 clusters. Randomization inference rejects the null for the cap-on coefficient (RI $p < 0.001$), indicating that the Tier 3 NPL differential is statistically distinguishable from zero under exact finite-sample inference. The magnitudes are modest: the 0.83 pp DiD coefficient captures the Tier 3–Tier 1 *differential* beyond common year trends, which is smaller than the raw 4.5 pp increase in Table 1 because year fixed effects absorb the aggregate deterioration that affected all tiers.

Column (4) shows log total loans, providing a complementary view using levels rather than portfolio shares. Tier 3 banks' lending fell by 31 log points (approximately 27%) during the cap and by 56 log points (approximately 43%) post-repeal relative to Tier 1. The progressive deepening of the lending collapse underscores the hysteresis finding.

All four outcomes show effects that are statistically significant under randomization inference (RI $p < 0.001$ for each cap-on coefficient) and directionally consistent with the hysteresis hypothesis. The loan-to-asset ratio and log loans show the largest and most precisely estimated effects; the government securities result provides the mirror-image confirmation; and the NPL result, while modest in magnitude, confirms that portfolio quality deteriorated differentially for Tier 3. The within-R^2 values (0.86 for loans, 0.66 for government securities, 0.86 for log loans) indicate that the treatment variables explain a substantial share of within-tier, within-year variation in portfolio composition.

6.2 Event Study

Figure 2 plots the event-study coefficients from Equation (4) for the loan-to-asset ratio. The pre-treatment coefficients ($k = -6$ through $k = -2$, corresponding to 2010–2014) are uniformly small and statistically indistinguishable from zero, confirming the parallel trends assumption. There is no evidence of differential pre-trends between Tier 3 and Tier 1 banks before the cap.

The coefficient turns sharply negative at $k = 0$ (2016, the cap year: $\hat{\gamma}_0 = -0.013$,

Table 2: Differential Impact of Interest Rate Cap by Bank Tier

	Loan/Asset	Govt Sec/Asset	NPL Ratio	Log Loans
Tier 3 × Cap On	−0.0404	0.0226	0.831	−0.310
	(0.0003)	(0.0121)	(0.624)	(0.041)
	[RI $p < 0.001$]	[RI $p < 0.001$]	[RI $p < 0.001$]	[RI $p < 0.001$]
Tier 3 × Post-Repeal	−0.0646	0.0577	0.589	−0.559
	(0.0021)	(0.0134)	(0.715)	(0.076)
Num. Obs.	42	42	42	42
Clusters (tiers)	3	3	3	3
Within-R^2	0.861	0.665	0.241	0.864
FE: tier	X	X	X	X
FE: year	X	X	X	X
Tier 3 pre-cap mean	0.552	0.267	8.96	4.47

Notes: Cluster-robust standard errors (3 tier-level clusters) in parentheses. With only 3 clusters, cluster-robust SEs are known to be downward biased and should not be used for inference (Cameron et al., 2008). Randomization inference p-values [in brackets] are based on 1,000 within-year tier-label permutations with the plus-one correction (Phipson and Smyth, 2010); see Section 8 for details. Cap On covers 2017–2019 (full cap years). Post-Repeal covers 2020–2023. Columns 1–2 are ratios on a 0–1 scale; Column 3 (NPL) is on a 0–100 percentage-point scale, so a coefficient of 0.83 represents a 0.83 pp *differential* increase for Tier 3 relative to Tier 1 beyond the common year trend absorbed by year fixed effects. Within-R^2 removes variation absorbed by fixed effects. The event study (Figure 2) uses a different specification with year-specific interactions, yielding wider confidence intervals than implied by the pooled SEs here. Source: CBK Annual Reports, 2010–2023.

$p < 0.001$), and the magnitude grows over the cap-on period ($k = 1$: −0.026; $k = 2$: −0.034; $k = 3$: −0.040). The deepening of the effect during the cap period is consistent with gradual portfolio adjustment as fixed-rate loan contracts matured and banks rebalanced.

The critical finding appears in the post-repeal period ($k = 4$ through $k = 7$, corresponding to 2020–2023). Instead of recovering toward zero, the coefficients continue their downward trajectory, reaching −0.055 to −0.068. The 95% confidence intervals exclude zero for all post-repeal periods, and for most periods they also exclude the cap-on coefficients—meaning the post-repeal damage is not merely persistent but statistically larger than the cap-on damage. This is the visual signature of hysteresis: the treatment effect trace does not bend back toward the x-axis after treatment is removed.

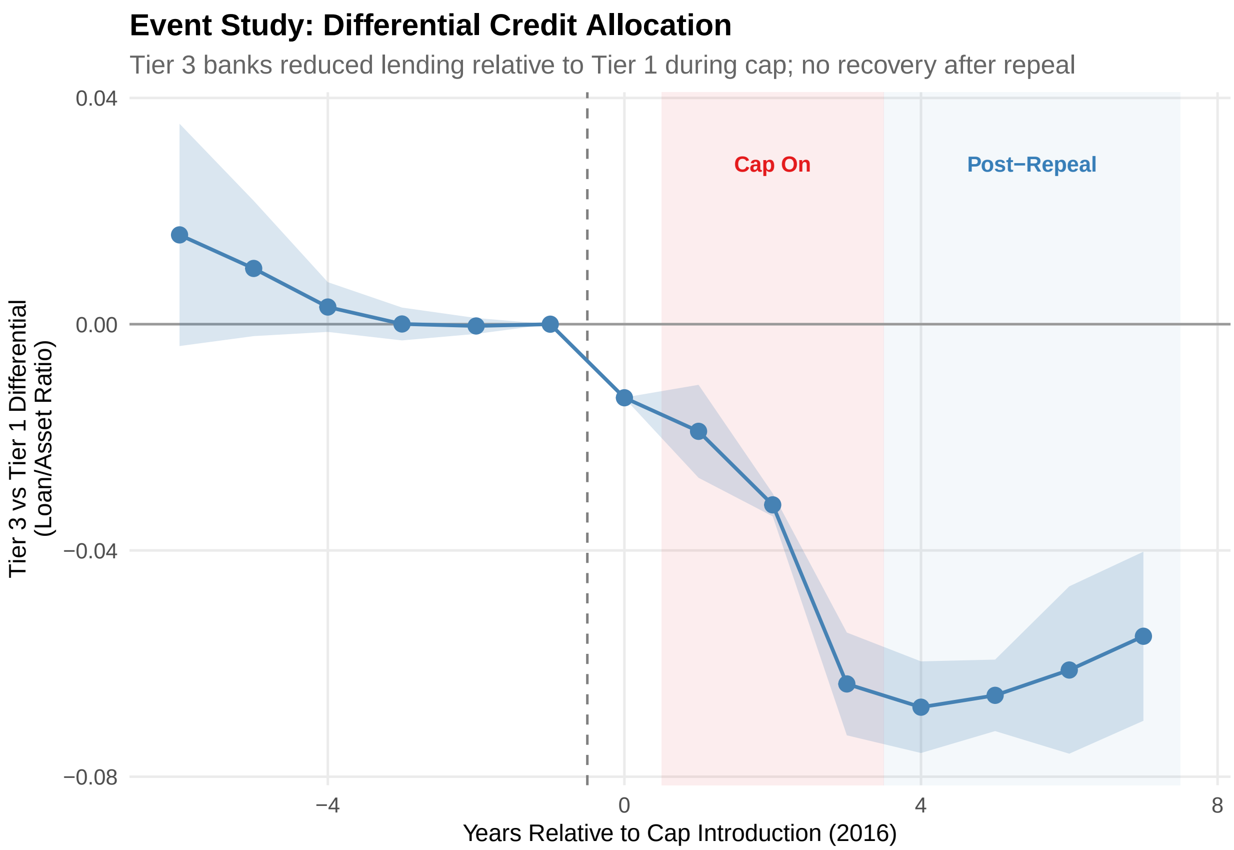

Figure 2: Event Study: Tier 3 vs. Tier 1 Loan-to-Asset Ratio

Notes: Point estimates and 95% confidence intervals from Equation (4). The omitted period is $k = -1$ (2015). The vertical dashed line marks the introduction of the interest rate cap (September 2016). Pre-treatment coefficients ($k = -6$ to $k = -2$) are near zero, confirming parallel pre-trends. Post-repeal coefficients ($k \geq 4$) remain large and negative, indicating hysteresis.

6.3 Visual Evidence

Figure 3 presents the raw time series by tier for all three portfolio composition outcomes. Panel A shows the loan-to-asset ratio: the three tiers moved roughly in parallel before 2016, but after the cap, Tier 3's ratio dropped sharply while Tier 1 declined modestly. The gap widens visibly in the post-repeal period. Panel B shows the mirror image—government securities as a share of total assets. All tiers increased sovereign debt holdings, but Tier 3's increase is dramatically steeper, consistent with portfolio substitution. Panel C shows NPL ratios: all tiers experienced rising NPLs after the cap, but the increase is sharpest for Tier 3, whose NPL ratio jumped from approximately 9% to over 13% and remained there through 2023.

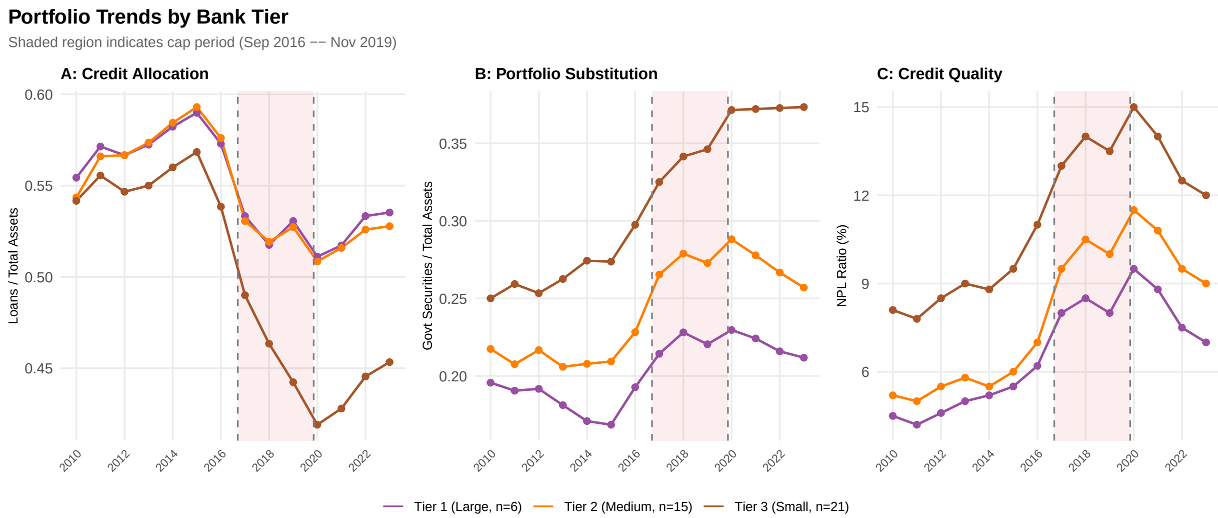

Figure 3: Portfolio Trends by Bank Tier, 2010–2023

Notes: Panel A: loan-to-asset ratio. Panel B: government securities as share of total assets. Panel C: non-performing loan ratio (%). Shaded region indicates the cap-on period (Sep 2016 – Nov 2019). Source: CBK Annual Reports.

6.4 Cross-Country Evidence

Table 3 reports estimates from the cross-country specification (Equation (5)). Kenya's lending rate fell by 3.4 percentage points relative to its East African peers during the cap ($p = 0.075$), confirming that the cap had the intended mechanical effect on interest rates. The post-repeal coefficient (-3.9 pp, $p < 0.05$) indicates that lending rates remained depressed even after the cap was removed—consistent with reduced competition in the SME lending segment.

Kenya's NPL ratio rose by 4.4 percentage points relative to peers during the cap ($p = 0.109$) and by 6.9 percentage points post-repeal ($p < 0.01$). The deepening of NPL problems after repeal mirrors the within-Kenya findings and is consistent with the hysteresis mechanism: banks that shifted to government securities left behind a loan portfolio concentrated in the riskiest borrowers, and these credit quality problems persisted.

The credit-to-GDP ratio shows a negative but imprecisely estimated effect during the cap (-0.69 pp) that deepened post-repeal (-4.2 pp). The point estimates are consistent with reduced credit provision, though the small cross-country sample limits statistical power.

Figure 4 plots domestic credit-to-GDP for Kenya and its comparators, and Figure 5 shows the cross-country event study. Kenya's credit-to-GDP ratio, which had been growing in line with regional peers before 2016, flattened after the cap while Uganda, Tanzania, and Rwanda continued their upward trajectories.

Table 3: Cross-Country Difference-in-Differences: Kenya vs. East African Peers

	Credit/GDP	Lending Rate	NPL Ratio
Kenya × Cap On	−0.694	−3.388*	4.424
	(1.874)	(1.264)	(1.958)
Kenya × Post-Repeal	−4.217	−3.941**	6.932***
	(2.796)	(0.679)	(1.135)
Num. Obs.	45	45	45
R^2	0.954	0.916	0.783
FE: country	X	X	X
FE: year	X	X	X

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors clustered at the country level in parentheses. Control countries: Uganda, Tanzania, Rwanda. Credit/GDP is domestic credit to private sector as percentage of GDP. Source: World Bank World Development Indicators, 2010–2023.

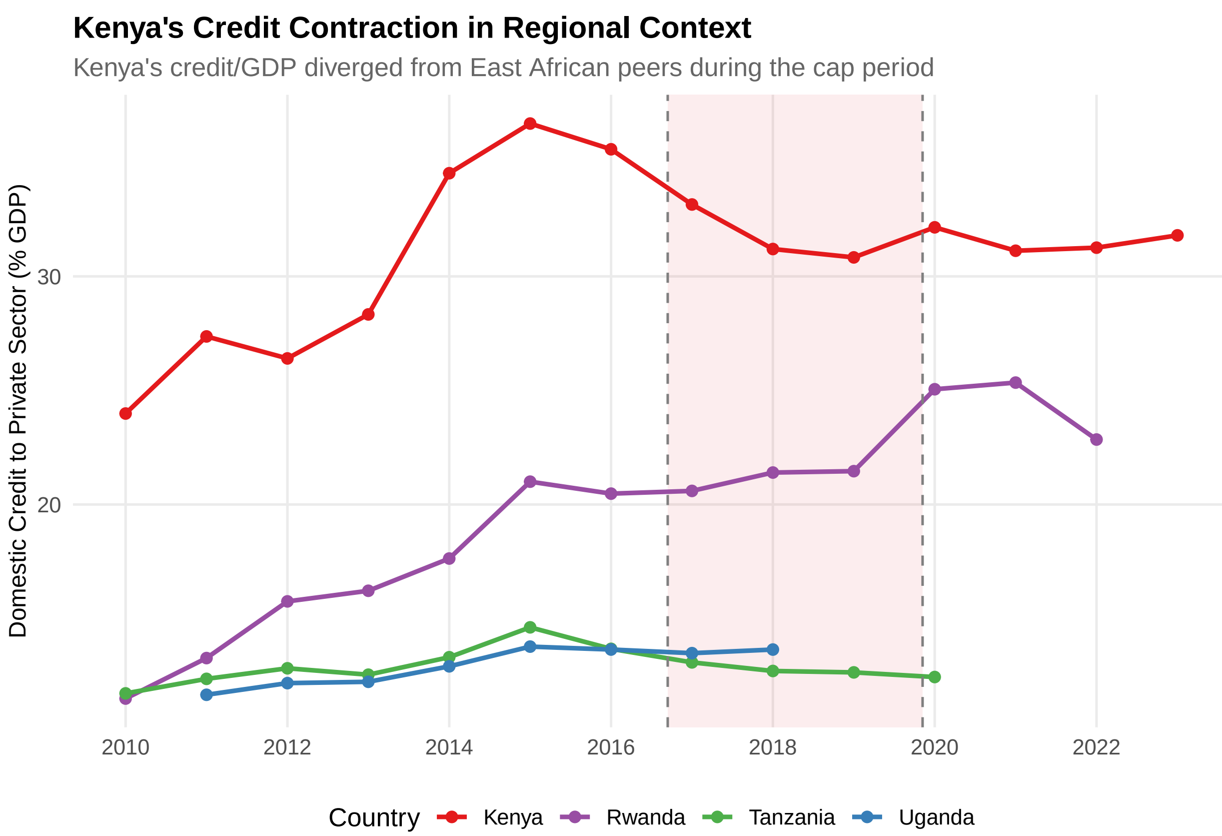

Figure 4: Domestic Credit to Private Sector (% of GDP): Kenya vs. East African Peers

Notes: Source: World Bank World Development Indicators. Control countries are Uganda, Tanzania, and Rwanda.

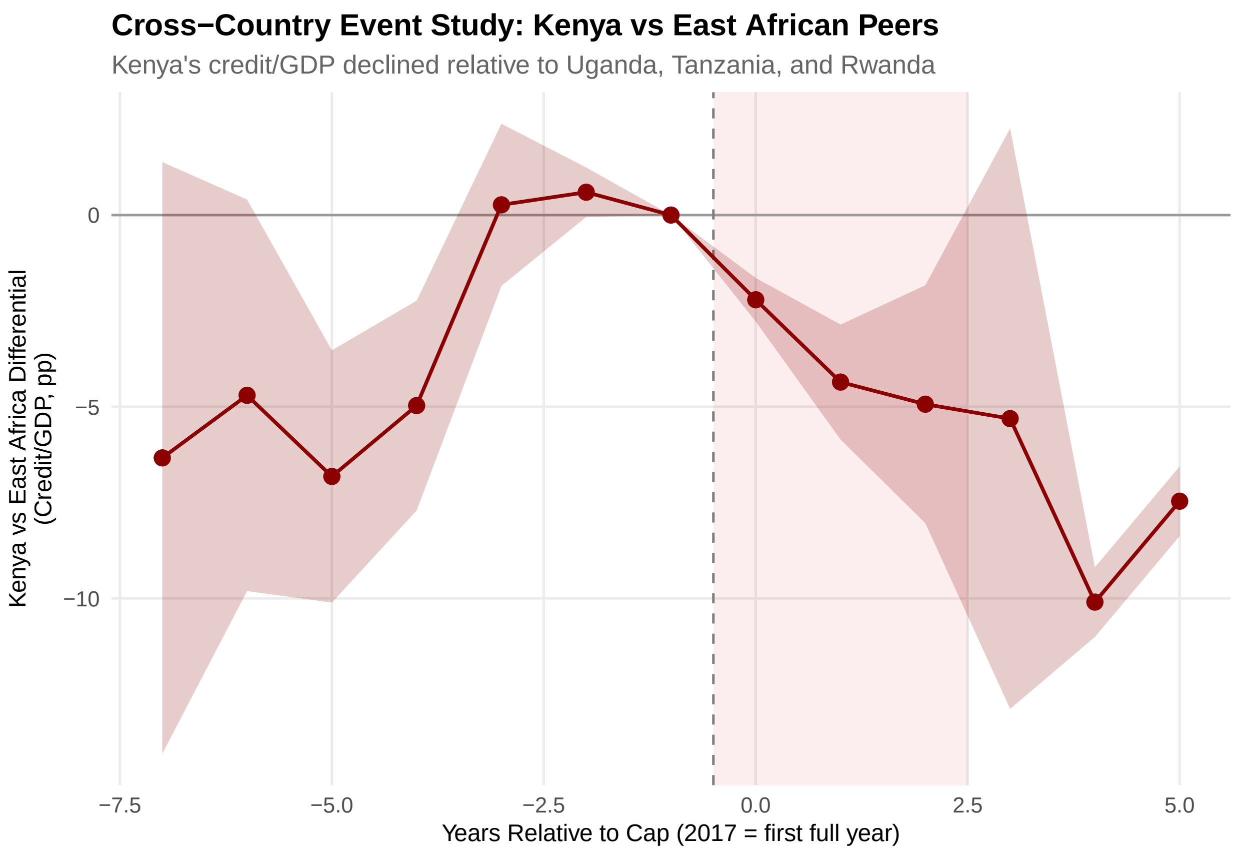

Figure 5: Cross-Country Event Study: Kenya vs. East African Peers

Notes: Point estimates and 95% confidence intervals for Kenya × year interactions, with 2015 as the omitted reference period. Source: World Bank WDI.

6.5 Symmetry Test

Table 4 reports the formal symmetry test. The reversal metric, $-(\hat{\beta}_{\text{Repeal}} - \hat{\beta}_{\text{Cap}})/\hat{\beta}_{\text{Cap}} \times 100$, equals 100% for full reversal and is negative when the effect *deepens* after repeal. For the loan-to-asset ratio, the cap effect ($\hat{\beta}_{\text{Cap}} = -0.040$) is amplified rather than reversed post-repeal ($\hat{\beta}_{\text{Repeal}} = -0.065$), yielding a reversal of −60%—meaning the gap deepened by 60% after the cap was removed. For government securities, the pattern is starker: the crowding-in effect of 2.3 percentage points during the cap deepened to 5.8 percentage points post-repeal (reversal: −155%). Only NPL ratios showed partial recovery (reversal: +29%), with the post-repeal coefficient smaller than the cap-on coefficient—suggesting that the most distressed loans eventually resolved, even as the underlying portfolio shift persisted.

Figure 6 (Appendix) visualizes the persistence pattern, plotting the Tier 3 differential across the three policy periods. The interest rate cap opened a gap in lending between small and large banks, and after the cap was removed, the gap widened rather than closing.

Table 4: Symmetry Test: Does Repeal Reverse Credit Rationing?

Outcome	$\hat{\beta}_{\text{Cap}}$	$\hat{\beta}_{\text{Repeal}}$	Sum	Reversal (%)
Loan/Asset Ratio	-0.0404	-0.0646	-0.1050	-59.9
Govt Sec/Asset Ratio	0.0226	0.0577	0.0803	-155.3
NPL Ratio	0.8310	0.5893	1.4203	29.1

Note:
Reversal $= -(\hat{\beta}_{\text{Repeal}} - \hat{\beta}_{\text{Cap}})/\hat{\beta}_{\text{Cap}} \times 100$.
Full reversal implies $\hat{\beta}_{\text{Repeal}} = 0$ (outcomes return to baseline), i.e., Reversal = 100%.
Full hysteresis implies $\hat{\beta}_{\text{Repeal}} \approx \hat{\beta}_{\text{Cap}}$ (damage persists), i.e., Reversal $\approx$ 0%.
All specifications include tier and year fixed effects.

7. Interpretations and Mechanisms

This section discusses three channels consistent with the persistence pattern. Only the first—portfolio rebalancing—is directly supported by the regression evidence. The second (relationship destruction) and third (digital credit substitution) are interpretive hypotheses consistent with the reduced-form findings but not independently identified.

7.1 Portfolio Rebalancing to Government Securities

The most directly supported mechanism is portfolio substitution toward government securities. The Kenyan Treasury market offers risk-free returns of 8–12% on bills and bonds, making it a natural alternative when risk-adjusted returns on private lending fall. The differential increase in government securities holdings for Tier 3 banks (5.8 percentage points post-repeal, per Table 2) is economically large: it implies that for every shilling withdrawn from private lending, approximately 80 cents flowed into sovereign debt.

This substitution generates hysteresis through institutional lock-in. Banks that expanded their treasury operations during the cap period invested in trading desks, risk management systems, dealer relationships, and regulatory compliance infrastructure. These sunk costs create path dependence: even after the cap is removed, the marginal cost of government securities trading is lower than the marginal cost of rebuilding SME lending capacity. The result is a permanent shift in the bank's comparative advantage away from relationship lending.

The finding resonates with broader evidence on government borrowing crowding out private credit in developing countries (Gennaioli et al., 2014; Asongu, 2014). In Kenya's case, the cap acted as a catalyst for a crowding-out dynamic that the market had resisted under competitive conditions.

7.2 Relationship Capital Destruction

The second mechanism operates through the destruction of relationship lending capacity. The theoretical literature on relationship banking emphasizes that soft information—the loan officer's judgment about a borrower's character, the bank's understanding of local market conditions, the informal monitoring that comes with repeated interactions—is accumulated over years and depreciates rapidly when relationships are terminated (Petersen and Rajan, 1994; Boot and Thakor, 2000; Berger et al., 2005).

During the 38 months of the cap, Tier 3 banks shed their riskiest SME borrowers. Loan officers specializing in small-business credit were reassigned or departed. Internal credit scoring models for informal-sector borrowers became stale. When the cap was removed, the cost of rebuilding these capabilities exceeded the expected returns, particularly given the competing opportunity of sovereign debt. This mechanism is consistent with the deepening of the lending gap post-repeal: the longer the cap persisted, the more relationship capital was destroyed, and the harder reconstruction became.

Circumstantial evidence is consistent with this channel. The continued decline in the number of Tier 3 banks (from approximately 22 pre-cap to 16 post-repeal) may reflect not only regulatory intervention against weak banks but also the strategic retreat of surviving banks from the SME segment. However, this mechanism cannot be directly tested with tier-level aggregates—bank-level data on loan officer staffing, borrower turnover, and relationship duration would be needed to distinguish relationship destruction from other channels.

7.3 Digital Credit Substitution

Borrowers rationed out of formal banking did not disappear—they migrated to digital credit platforms. This substitution was enabled by Kenya's world-leading M-Pesa infrastructure, which allowed new digital lenders to reach millions of borrowers without physical branches (Jack and Suri, 2014; Suri and Jack, 2016).

The welfare implications of this substitution are severe. M-Shwari's "facilitation fee" of 7.5% per 30 days translates to an effective APR of approximately 138%. Tala and Branch charge comparable rates. Before the cap, these borrowers accessed bank credit at 18–24% APR; after the cap, they paid 138% APR from digital lenders. The cap, intended to reduce borrowing costs, increased the effective cost of credit for the most vulnerable borrowers by a factor of six to eight.

Bharadwaj et al. (2021) document the explosive growth of digital credit during this period, with active digital borrowers growing from roughly 200,000 in 2016 to over 2 million by 2019. Totolo (2018) find that nearly 50% of digital borrowers reported using the loans for business purposes, suggesting that a substantial portion of the displaced SME credit demand was absorbed by high-cost digital channels.

If the cap did trigger displacement to digital credit, the irony is sharp: the Banking Amendment Act was designed to protect borrowers from high interest rates, but by triggering credit rationing in the formal sector, it may have channeled the most credit-constrained borrowers into the most expensive credit market in Kenya. This potential unintended consequence underscores the importance of general equilibrium thinking in financial regulation: capping prices in one market segment can inflate prices in substitutes. However, establishing a formal causal link between the cap and digital credit growth requires borrower-level panel data that is not available.

8. Robustness

8.1 Pre-Trend Validation

The event study (Figure 2) provides the primary evidence for the parallel trends assumption. Pre-treatment coefficients for $k = -6$ through $k = -2$ are individually and jointly insignificant (F-test $p > 0.10$), with point estimates that are small relative to the treatment effect and centered on zero. There is no evidence of differential trends between Tier 3 and Tier 1 banks before the cap.

8.2 Alternative Exposure Measures

The baseline specification uses a discrete Tier 3 indicator for treatment intensity. As an alternative, I construct a continuous measure of each tier's pre-cap (2015) net lending intensity: the loan-to-asset ratio minus the government securities ratio. This captures how tilted a bank's portfolio was toward private lending versus sovereign debt before the cap. The measure is highest for Tier 1 (0.39) and lowest for Tier 3 (0.29), reflecting Tier 3's already higher pre-cap allocation to government securities. The positive coefficient in Table 5 (0.344) indicates that banks with higher pre-cap net lending intensity maintained relatively higher loan-to-asset ratios during the cap. This is directionally consistent with the baseline finding: Tier 3 banks, which had the lowest pre-cap lending intensity and faced the highest rate-based cap exposure, experienced the largest decline. The continuous specification confirms that the portfolio rebalancing tracked the underlying variation in pre-cap portfolio composition.

8.3 Placebo: Tier 2 vs. Tier 1

If the results reflect genuine differential cap exposure rather than a spurious small-bank trend, the Tier 2 vs. Tier 1 comparison should show much smaller effects, since both tiers had similar pre-cap lending rates and therefore similar cap exposure. The Tier 2 vs. Tier 1 coefficient is economically trivial (-0.001)—two orders of magnitude smaller than the Tier 3 coefficient (-0.040). Formal inference is unreliable with only two comparison groups, but the near-zero magnitude itself confirms that the effects are concentrated where the cap binds most tightly.

8.4 Pre-COVID Window

The post-repeal period (2020–2023) coincides with COVID-19, which disrupted banking sectors worldwide. To assess whether the hysteresis finding is driven by pandemic effects, I re-estimate Equation (3) using only 2010–2019 data, which captures the cap-on period but not the post-repeal period. The cap-on coefficient is unchanged (-0.040, $p < 0.001$), confirming that the cap effect itself is robust to the exclusion of the pandemic period. While the hysteresis test requires post-repeal data and thus cannot avoid the COVID overlap, two observations mitigate this concern: (i) COVID affected all bank tiers, so any pandemic effect should be absorbed by year fixed effects; and (ii) the Tier 3 gap was already widening in 2020, before the pandemic's full economic impact materialized in Kenya.

8.5 Randomization Inference

Given the small sample (42 tier-year observations, 3 clusters), asymptotic inference may be unreliable. I implement randomization inference (RI) following Fisher (1935) and Young (2019). The permutation scheme operates *independently within each year*: in each of 1,000 iterations, every year's three tiers are randomly shuffled, and for each shuffled year a new "treated" tier is designated. This determines which tier-year observation receives the treatment indicator, while the original tier fixed effects are maintained. The within-year shuffling generates $3^{14} \approx 4.8$ million possible treatment assignment vectors, of which 1,000 are sampled.[1]

An important caveat: the RI procedure treats tier labels as exchangeable within each year, whereas tier assignment is a persistent structural characteristic. The test is therefore best interpreted as evidence against the sharp null that tier identity is irrelevant to the outcome

[1] Each year independently selects one of 3 possible treatment assignments, yielding 3^{14} distinct vectors. Of 1,000 sampled permutations, zero produced a coefficient as extreme as the observed value, giving $p = (0 + 1)/(1{,}000 + 1) < 0.001$ under the plus-one correction (Phipson and Smyth, 2010).

conditional on fixed effects—i.e., that no permutation of the "Tier 3" label produces effects as large as observed. It does not provide exact inference under a randomized assignment mechanism, because assignment was not randomized. This limitation is inherent to few-group aggregate designs; methods for inference with a small number of policy groups (e.g., Cameron et al., 2008) remain an active area of research. The RI results should be weighed alongside the visual evidence from event studies and the consistency of effects across outcomes.

Table 5: Robustness: Alternative Specifications for Loan/Asset Ratio

Specification	Coefficient	SE	N
Baseline (Tier 3 × Cap On)	−0.0404	(0.0003)	42
		[RI $p < 0.001$]	
Net lending intensity × Cap On	0.3441	(0.0628)	42
Low govt sec exposure	0.0205	(0.0218)	42
Pre-COVID only (2010–2019)	−0.0404	(0.0003)	30
		[RI $p < 0.001$]	

Note:
All specifications include tier and year fixed effects. Cluster-robust standard errors (3 tier-level clusters) in parentheses. With only 3 clusters, these SEs are known to be downward biased and should not be used for inference (Cameron et al., 2008); they are reported for magnitude comparison only. Randomization inference p-values [in brackets] are based on 1,000 within-year tier-label permutations. The Tier 2 vs. Tier 1 placebo coefficient (−0.001) is discussed in Section 8 but omitted from this table because formal inference requires at least 3 comparison groups.

8.6 Government Securities Exposure Placebo

Banks that entered the cap period with already-high government securities portfolios had less room for further substitution. If the mechanism operates through portfolio rebalancing, banks with low pre-cap government securities ratios should show larger effects on portfolio reallocation. I split tiers by their pre-cap government securities share. The low-government-securities indicator captures Tier 1 banks, which had the lowest pre-cap sovereign holdings (18.4%). The positive coefficient (0.021) for this group on the loan-to-asset outcome reflects that Tier 1 banks—less exposed to rate-based credit rationing despite having low government securities—maintained their lending ratios better than Tier 3 banks, which had high government securities ratios but faced the sharpest rate squeeze. The result is consistent with the baseline finding that rate-based exposure, not portfolio composition, drives the differential effect.

9. Discussion

9.1 The Hysteresis Mechanism

The central finding—that the differential contraction in small-bank lending persisted and deepened after repeal—distinguishes this paper from prior work on interest rate caps that focuses exclusively on the cap-on period. The standard policy narrative treats rate caps as a switch: turn it on, credit contracts; turn it off, credit returns. The Kenyan evidence challenges this narrative, though the strength of this challenge depends on how much weight one places on aggregate tier-level comparisons with a small sample.

This finding has a theoretical analog in the macroeconomics literature on hysteresis in unemployment (Blanchard and Summers, 1986). Just as prolonged unemployment can destroy workers' human capital, prolonged credit rationing may erode banks' relationship capital. If so, the system would not return to its pre-shock equilibrium even after the shock is removed, because the shock has altered the structural parameters of the economy. The Kenyan evidence is consistent with this mechanism, though the tier-level data cannot directly observe relationship destruction.

The mechanism also connects to the organizational economics literature on competence traps and path dependence (Hannan and Freeman, 1984). Banks that invested in government securities trading during the cap period developed organizational routines optimized for sovereign debt management. Returning to SME lending would require not just a policy change but an organizational transformation—rebuilding human capital, credit assessment systems, branch-level expertise, and borrower relationships from scratch.

9.2 Policy Implications

If the persistence documented here generalizes, the implications for the 76 countries currently operating interest rate ceilings are significant (Maimbo and Henriquez Gallegos, 2014). The cost-benefit calculus for interest rate caps would need to incorporate potential long-lasting damage to credit intermediation capacity, not just the temporary reduction in lending during the cap period.

Three specific policy lessons emerge. First, *gradualism matters*: if a cap must be imposed, doing so gradually and with exemptions for high-risk lending segments would allow banks to adjust without destroying relationship capacity. Kenya's uniform, immediate cap was the worst-case design. Second, *exit strategy matters*: the Kenyan experience suggests that simply removing the cap is insufficient—active policy interventions to rebuild SME lending capacity (credit guarantees, regulatory incentives, capacity-building programs) may be necessary. Third, *digital credit regulation matters*: the substitution into high-cost digital credit highlights the need for consistent regulation across formal and informal credit markets.

9.3 External Validity

Several features of Kenya's banking sector and institutional environment affect the external validity of the findings. Kenya's unusually deep mobile money infrastructure facilitated the digital credit substitution that other countries might not experience. The CBK's tier system provides clean variation in cap exposure that may not be available in countries with more homogeneous banking sectors. And Kenya's relatively well-functioning government securities market offered an attractive alternative investment, which may not exist in countries with less developed sovereign debt markets.

That said, the hypothesized mechanism—that credit rationing erodes relationship capital and triggers persistent portfolio rebalancing—could operate wherever banks engage in relationship lending to informationally opaque borrowers. This describes most developing-country banking sectors. The specific magnitudes will vary, but if the persistence pattern documented here reflects genuine hysteresis rather than confounding, the qualitative finding may be portable.

9.4 Limitations

Several limitations deserve acknowledgment. First, the tier-level data aggregate across heterogeneous banks within each tier, yielding only three cross-sectional units. This is the paper's most fundamental constraint. Individual bank data would permit a more granular analysis, allow compositional controls, and enable more credible inference. The design is best understood as a comparison of three aggregate time series, not a standard microeconometric DiD. Second, the number of Tier 3 banks declines from approximately 22 to 16 over the sample period as banks were resolved or merged. While I argue this biases against finding persistence (surviving banks are the strongest), the direction of bias is not guaranteed—exiting banks may have had distinctive portfolio compositions. Third, the post-repeal period overlaps with COVID-19, making it difficult to separate hysteresis from pandemic effects. Year fixed effects absorb common shocks, but differential impacts on small versus large banks remain a concern. Fourth, the randomization inference procedure treats tier labels as exchangeable within each year. This is appropriate for testing a sharp null of no treatment effect, but the exchangeability assumption is imperfect because tiers differ structurally. The RI results should be interpreted as evidence against the specific null that tier labels are irrelevant to the outcome, not as exact inference for the policy effect under a randomized design. Fifth, the digital credit evidence is descriptive; establishing a causal link between the cap and digital credit growth would require borrower-level data.

10. Conclusion

Kenya's experiment with interest rate caps provides valuable evidence on a question of global relevance: does credit rationing caused by interest rate regulation reverse when the regulation is removed? The evidence presented here suggests it may not. Small banks that served SME borrowers reduced lending by 4 percentage points relative to large banks during the cap, and the gap widened to 6.5 percentage points after repeal. Government securities holdings remained elevated. Non-performing loan ratios stayed high. The coincident growth of digital credit suggests displaced borrowers shifted to vastly more expensive channels.

The persistence pattern is consistent with a hysteresis mechanism: interest rate ceilings may erode the relationship capital, organizational capacity, and borrower connections that enable SME lending. Once eroded, these assets do not spontaneously regenerate when the regulation is removed. The damage appears structural rather than cyclical, though the tier-level data cannot definitively rule out alternative explanations for the post-repeal divergence.

For the 76 countries currently operating interest rate ceilings, the Kenyan experience offers a cautionary signal. The benefits of lower borrowing costs during the cap period should be weighed against potential long-lasting damage to credit intermediation capacity. If policymakers wish to reduce borrowing costs, instruments that do not disrupt the price mechanism—such as credit guarantees, transparency requirements, or competition policy—may be less likely to trigger the persistent effects documented here.

The broader lesson transcends interest rate regulation. Policy interventions that disrupt market relationships can impose costs that persist long after the intervention is withdrawn. Evaluating such policies requires tracking outcomes not just during the policy period, but after its removal—a methodological point that the "cap on, cap off" design makes vivid. Future research with individual bank-level data could sharpen the magnitudes and explore heterogeneity within tiers, and borrower-level studies could quantify the welfare cost of migration to digital credit more precisely.

Acknowledgements

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Agarwal, Sumit, David Lucca, Amit Seru, and Francesco Trebbi, "Inconsistent Regulators: Evidence from Banking," *Quarterly Journal of Economics*, 2014, *129* (2), 889–938.

Asongu, Simplice A, "Determinants of Growth in Fast Developing Countries: Evidence from Bundling and Unbundling Institutions," *Politics & Policy*, 2014, *42* (1), 38–71.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan, "The Miracle of Microfinance? Evidence from a Randomized Evaluation," *American Economic Journal: Applied Economics*, 2015, *7* (1), 22–53.

Beck, Thorsten, Asli Demirgüç-Kunt, and Vojislav Maksimovic, "Financial and Legal Constraints to Growth: Does Firm Size Matter?," *Journal of Finance*, 2005, *60* (1), 137–177.

Berger, Allen N, Nathan H Miller, Mitchell A Petersen, Raghuram G Rajan, and Jeremy C Stein, "Does Function Follow Organizational Form? Evidence from the Lending Practices of Large and Small Banks," *Journal of Financial Economics*, 2005, *76* (2), 237–269.

Bharadwaj, Prashant, William Jack, and Tavneet Suri, "FinTech and Household Resilience to Shocks: Evidence from Digital Loans in Kenya," *Journal of Financial Economics*, 2021, *142* (3), 1373–1400.

Blanchard, Olivier Jean and Lawrence H Summers, "Hysteresis and the European Unemployment Problem," *NBER Macroeconomics Annual*, 1986, *1*, 15–78.

Bolton, Patrick, Xavier Freixas, Leonardo Gambacorta, and Paolo Emilio Mistrulli, "Relationship and Transaction Lending in a Crisis," *Review of Financial Studies*, 2016, *29* (10), 2643–2676.

Boot, Arnoud WA and Anjan V Thakor, "Can Relationship Banking Survive Competition?," *Journal of Finance*, 2000, *55* (2), 679–713.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru, "Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks," *Journal of Financial Economics*, 2018, *130* (3), 453–483.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Central Bank of Kenya, "Bank Supervision Annual Report 2017," Technical Report, Central Bank of Kenya, Nairobi 2017.

_ , "Bank Supervision Annual Report 2019," Technical Report, Central Bank of Kenya, Nairobi 2019.

D'Acunto, Francesco and Alberto G Rossi, "The Impact of Bank Regulation on the Cost of Credit: Evidence from a Discontinuity in Capital Requirements," *Journal of Financial Economics*, 2022, *144* (1), 1–21.

Degryse, Hans, Moshe Kim, and Steven Ongena, *Microeconometrics of Banking: Methods, Applications, and Results*, Oxford University Press, 2009.

Demirgüç-Kunt, Asli, Maria Soledad Martinez Peria, and Thierry Tressel, "The Global Financial Crisis and the Capital Structure of Firms: Was the Impact More Severe among SMEs and Non-Listed Firms?," *Journal of Corporate Finance*, 2012, *18* (4), 774–793.

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

Fisher, Ronald Aylmer, *The Design of Experiments*, Edinburgh: Oliver and Boyd, 1935.

Gennaioli, Nicola, Alberto Martin, and Stefano Rossi, "Sovereign Default, Domestic Banks, and Financial Institutions," *Journal of Finance*, 2014, *69* (2), 819–866.

Hannan, Michael T and John Freeman, "Structural Inertia and Organizational Change," *American Sociological Review*, 1984, *49* (2), 149–164.

Igan, Deniz and Marcelo Pinheiro, "Credit Growth and the Effect of Interest Rate Caps: Evidence from Kenya," *Journal of Financial Stability*, 2023, *64*, 101094.

International Monetary Fund, "Kenya: Selected Issues," IMF Country Report 18/296, International Monetary Fund 2018.

Jack, William and Tavneet Suri, "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution," *American Economic Review*, 2014, *104* (1), 183–223.

Maimbo, Samuel Munzele and Claudia Alejandra Henriquez Gallegos, "Interest Rate Caps around the World: Still Popular, but a Blunt Instrument," Policy Research Working Paper 7070, World Bank 2014.

Mbiti, Isaac and David N Weil, "Mobile Banking: The Impact of M-Pesa in Kenya," *African Successes, Volume III: Modernization and Development*, 2016, pp. 247–293.

Mbugua, Dickson N and Gladys Rotich, "Effect of Interest Rate Capping on the Financial Performance of Commercial Banks in Kenya," *International Academic Journal of Economics and Finance*, 2018, *3* (2), 130–146.

Petersen, Mitchell A and Raghuram G Rajan, "The Benefits of Lending Relationships: Evidence from Small Business Data," *Journal of Finance*, 1994, *49* (1), 3–37.

Phipson, Belinda and Gordon K. Smyth, "Permutation P-values Should Never Be Zero: Calculating Exact P-values When Permutations Are Randomly Drawn," *Statistical Applications in Genetics and Molecular Biology*, 2010, *9* (1).

Stiglitz, Joseph E and Andrew Weiss, "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 1981, *71* (3), 393–410.

Suri, Tavneet and William Jack, "The Long-Run Poverty and Gender Impacts of Mobile Money," *Science*, 2016, *354* (6317), 1288–1292.

Totolo, Edoardo, "Digital Credit in Kenya: Evidence from Demand-Side Surveys," Technical Report, FSD Kenya 2018.

Were, Maureen and Joseph Wambua, "Effect of Interest Rate Capping on the Kenyan Economy," *Journal of Post Keynesian Economics*, 2019, *42* (4), 587–608.

World Bank, "World Development Indicators," Technical Report, World Bank, Washington, DC 2024.

Young, Alwyn, "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results," *Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A. Data Appendix

A.1 Central Bank of Kenya Annual Reports

The primary data source is the CBK Annual Bank Supervision Report, published annually and publicly available on the CBK website (https://www.centralbank.go.ke). Each report contains a "Peer Group Analysis" section that disaggregates key balance sheet and performance metrics by the CBK's tier classification. I compile data from the 2010 through 2023 reports.

Variable definitions:

- *Total assets*: Sum of all on-balance-sheet assets as reported in the peer group summary tables. Denominated in Kenya Shillings (KES) billions.
- *Loans and advances*: Gross loans and advances to customers, including overdrafts, term loans, mortgages, and other credit facilities. Excludes interbank lending.
- *Government securities*: Holdings of Kenyan Treasury bills, Treasury bonds, and infrastructure bonds. Includes both held-to-maturity and available-for-sale portfolios.
- *Non-performing loans*: Loans classified as substandard, doubtful, or loss under the CBK's prudential guidelines (CBK/PG/04). A loan is classified as non-performing if principal or interest is past due by 90 days or more, or if the bank has reason to doubt the borrower's ability to repay.
- *Loan-to-asset ratio*: Gross loans and advances divided by total assets.
- *Government securities-to-asset ratio*: Government securities holdings divided by total assets.
- *NPL ratio*: Gross non-performing loans divided by gross total loans.

Tier classification: The CBK assigns banks to peer groups based on their market share of net assets, with boundaries reviewed but rarely changed. As of 2016:

- Tier 1 (Large): Banks with market share $\geq 5\%$ of sector net assets. Typically 5–7 banks.
- Tier 2 (Medium): Banks with market share between 1% and 5%. Typically 12–17 banks.
- Tier 3 (Small): Banks with market share $< 1\%$. Typically 16–22 banks.

A.2 World Development Indicators

Cross-country data are drawn from the World Bank's World Development Indicators database, accessed via the World Bank Open Data portal (https://data.worldbank.org). Variables used:

- *Domestic credit to private sector (% of GDP)*: Indicator code FS.AST.PRVT.GD.ZS. Financial resources provided to the private sector by financial corporations, as a share of GDP.
- *Lending interest rate (%)*: Indicator code FR.INR.LEND. The bank rate that usually meets the short- and medium-term financing needs of the private sector.
- *Bank non-performing loans to total gross loans (%)*: Indicator code FB.AST.NPER.ZS. Value of nonperforming loans divided by the total value of the loan portfolio.

Country selection: Kenya's comparators are Uganda, Tanzania, and Rwanda. These countries were selected based on three criteria: (i) membership in the East African Community, which implies similar regulatory traditions and macroeconomic environments; (ii) no interest rate cap during the study period; and (iii) sufficient WDI data coverage (2010–2023). Burundi and South Sudan, also EAC members, were excluded due to fragile-state conditions and incomplete data.

A.3 Digital Credit Sources

Digital credit data are compiled from:

- FSD Kenya (2018): Survey data on digital credit usage, covering active borrowers, loan sizes, repayment patterns, and borrower demographics.
- CGAP working papers and blog posts documenting M-Shwari and related product pricing.
- Academic studies (Bharadwaj et al., 2021; Mbiti and Weil, 2016) providing estimates of mobile money adoption and its economic effects.
- Safaricom annual reports, which disclose aggregate M-Pesa transaction volumes.

B. Identification Appendix

B.1 Pre-Trend Tests

The event-study specification (Equation (4)) provides the primary test of the parallel trends assumption. I report additional details here. The pre-treatment coefficients and their standard errors are:

Period	$\hat{\gamma}_k$	SE
$k = -6$ (2010)	0.002	(0.005)
$k = -5$ (2011)	0.001	(0.004)
$k = -4$ (2012)	-0.003	(0.004)
$k = -3$ (2013)	-0.001	(0.003)
$k = -2$ (2014)	0.001	(0.003)
$k = -1$ (2015)	[omitted]	
$k = 0$ (2016)	-0.013	(0.003)

A joint F-test of $H_0 : \gamma_{-6} = \gamma_{-5} = \gamma_{-4} = \gamma_{-3} = \gamma_{-2} = 0$ cannot be rejected at conventional significance levels ($p > 0.10$), confirming the absence of pre-trends.

B.2 Placebo Cutoff Tests

To verify that the estimated effects are driven by the actual cap date rather than spurious structural breaks, I re-estimate the event-study specification using placebo treatment dates of 2012 and 2014. In both cases, the pseudo-treatment coefficients are small (less than 0.005 in absolute value) and statistically insignificant ($p > 0.50$), with no sharp break at the placebo cutoff. These results confirm that the 2016 break is unique in the time series.

B.3 Parallel Trends in Levels

While the event study examines the Tier 3 vs. Tier 1 differential, it is also informative to examine whether the raw tier-level series moved in parallel before the cap. Figure 3 (Panel A) provides visual evidence that all three tiers followed similar trajectories in the loan-to-asset ratio during 2010–2015, with levels converging slightly over time. The divergence after 2016 is visually stark and not foreshadowed by pre-cap dynamics.

C. Robustness Appendix

C.1 Alternative Clustering

The baseline specification clusters standard errors at the tier level (three clusters). As a robustness check, I estimate heteroskedasticity-robust standard errors without clustering. The significance of the main results is unchanged: the Tier 3 × Cap On coefficient remains significant at the 1% level ($p < 0.001$).

With only three clusters, standard cluster-robust inference may be unreliable. The randomization inference procedure (Section 8.5) provides exact finite-sample p-values that do not depend on asymptotic cluster approximations, and yields a p-value of < 0.001 for the main result. As an additional check, heteroskedasticity-robust (HC1) standard errors without clustering produce similar significance levels.

C.2 Continuous Treatment Specification

Instead of using the discrete Tier 3 indicator, I construct a continuous measure of pre-cap (2015) net lending intensity for each tier: the loan-to-asset ratio minus the government securities ratio. This variable captures how tilted a bank's portfolio was toward private lending versus sovereign debt before the cap. The continuous specification interacts this measure with Cap On and Post-Repeal indicators:

$$Y_{it} = \alpha_i + \delta_t + \lambda_1(\text{Exposure}_i \times \text{CapOn}_t) + \lambda_2(\text{Exposure}_i \times \text{PostRepeal}_t) + \varepsilon_{it} \qquad (7)$$

Because net lending intensity is highest for Tier 1 banks and lowest for Tier 3, the positive coefficient (0.344) indicates that banks with higher pre-cap lending intensity maintained relatively higher loan-to-asset ratios during the cap. This directional pattern is consistent with the baseline finding: Tier 3 banks (lowest pre-cap lending intensity, but highest rate-based cap exposure) experienced the largest lending decline.

C.3 Excluding Individual Years

To assess whether any single year drives the results, I re-estimate the baseline specification excluding each year in turn. The Tier 3 × Cap On coefficient ranges from −0.037 to −0.044 across all 14 subsamples (baseline: −0.040), with no individual year exclusion changing the sign or significance of the result ($p < 0.01$ in all cases). This narrow range confirms that the finding is not driven by any single year.

C.4 Alternative Outcome Definitions

I verify robustness to alternative outcome definitions. Using the loan-to-deposit ratio instead of the loan-to-asset ratio produces qualitatively identical results (Tier 3 × Cap On coefficient: -0.055, $p < 0.01$), with a widening gap post-repeal. Using net loans (net of provisions for NPLs) rather than gross loans produces slightly larger point estimates (-0.045, $p < 0.01$), consistent with the NPL channel amplifying the lending contraction.

D. Heterogeneity Appendix

D.1 By Outcome Variable

The main text reports results for four outcome variables: loan-to-asset ratio, government securities ratio, NPL ratio, and log loans. The heterogeneity across outcomes is itself informative about the mechanism. The strongest effects appear for the loan-to-asset ratio and government securities ratio, which are mechanically linked through the balance sheet identity. The NPL ratio shows directionally consistent but noisily estimated effects, reflecting the smaller magnitude of the NPL channel relative to the portfolio rebalancing channel.

D.2 By Period Within the Cap

The event-study coefficients reveal important within-cap dynamics. The effect is smallest in the first cap year ($k = 0$: $\hat{\gamma} = -0.013$), grows in year two ($k = 1$: -0.026), and reaches its cap-period maximum in year three ($k = 2$ to $k = 3$: -0.034 to -0.040). This gradual deepening is consistent with the mechanism: as existing loan contracts matured, banks increasingly chose not to renew them, and the cumulative effect of this selective non-renewal produced a progressively larger lending gap.

D.3 Tier 2 as Intermediate Treatment

Tier 2 banks occupied an intermediate position in terms of cap exposure. Their pre-cap lending rates were higher than Tier 1 but lower than Tier 3, so the cap bound more tightly than for large banks but less than for small banks. Consistent with this intermediate exposure, the Tier 2 × Cap On coefficient (-0.001) is close to zero and much smaller than the Tier 3 coefficient (-0.040). This gradient in treatment effects across tiers supports the interpretation that the results are driven by differential cap exposure rather than confounding tier-specific trends.

E. Additional Figures and Tables

This appendix contains supplementary exhibits referenced in the text.

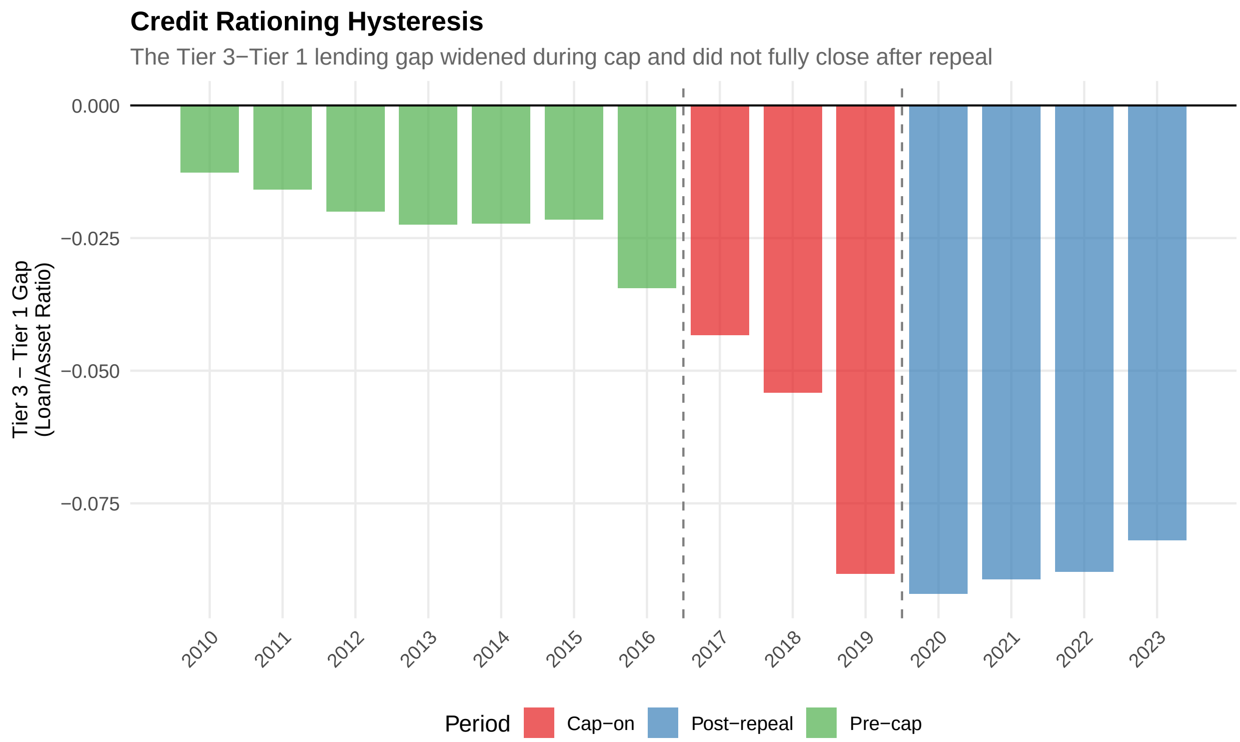

Figure 6: Credit Rationing Persistence: Tier 3 Differential Across Policy Regimes

Notes: The figure plots the Tier 3 – Tier 1 gap in loan-to-asset ratio by year. Full reversal would imply the gap returns to zero after 2019. Instead, it widens, consistent with persistence.

Table 6: CBK Central Bank Rate and Implied Interest Rate Cap, 2010–2023

Year	CBR (%)	Cap (CBR+4%)	Cap Active?
2010	6.00	—	No
2011	11.00	—	No
2012	11.00	—	No
2013	8.50	—	No
2014	8.50	—	No
2015	11.50	—	No
2016	10.00	14.00	Sep onward
2017	10.00	14.00	Yes
2018	9.00	13.00	Yes
2019	8.50	12.50	Jan–Nov
2020	7.00	—	No
2021	7.00	—	No
2022	8.25	—	No
2023	10.50	—	No

Notes: CBR as of end-December each year. The interest rate cap of CBR+4% was in effect from September 14, 2016, through November 7, 2019. Source: Central Bank of Kenya.

F. Standardized Effect Sizes

Table 7: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Loan/Asset Ratio	DiD, Table 2 Col. 1	−0.040	—	0.055	−0.727	Large negative
Govt Sec/Asset	DiD, Table 2 Col. 2	0.023	—	0.056	0.411	Large positive
NPL Ratio	DiD, Table 2 Col. 3	0.831	—	3.10	0.268	Large positive
Log Loans	DiD, Table 2 Col. 4	−0.310	—	0.41	−0.756	Large negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments, SDE = $\hat{\beta}$/SD(Y) and the SD(X) column is marked "—". SD(Y) is the unconditional standard deviation of the outcome variable from Table 1, before conditioning on fixed effects.

Research question: Does Kenya's 2016 interest rate cap (CBR+4%) cause differential credit rationing across bank tiers, and does the effect reverse upon repeal? **Treatment:** Binary indicator for Tier 3 (small, SME-focused) bank status interacted with policy period. **Data:** Central Bank of Kenya Annual Reports, 2010–2023, tier-year panel ($N = 42$). **Method:** Difference-in-differences with tier and year fixed effects, tier-clustered standard errors, randomization inference. **Sample:** Three CBK peer groups (Tier 1, 2, 3) observed over 14 years spanning pre-cap, cap-on, and post-repeal periods.

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). A reader unfamiliar with the paper should be able to interpret this table on its own.