Allocative Efficiency in Bank Internal Capital Markets*

Deniz Aydın — MIT

Ernest Liu — Princeton

Janis Skrastins — Utah

July 2026

Abstract

We study whether a bank's internal capital market allocates capital efficiently across branches. Using a large-scale randomized pricing experiment on approximately 50,000 loan applications, we identify the marginal return on capital for each branch group—the change in profit per dollar of capital allocated. We find that total lending is close to its optimal scale, but marginal returns are not equalized across branch groups, implying potential reallocation gains of 8.4% of profits. The bank's uniform pricing policy—charging the same spread across all branches conditional on credit score—impedes price-based reallocation.

Keywords: internal capital markets, capital allocation, uniform pricing, field experiment, banking

JEL: C93, D61, G21, G31, G51

*We thank Cevdet Akcay, Paul Carrillo, Philip Dybvig, Ali Hortacsu, David Thesmar, Pete Klenow, Jeremy Majerovitz, Antoinette Schoar, David Sraer, Adi Sunderam; our discussants Kim Fe Cremer, Rustom Irani, Martin Oehmke, Gyongyi Loranth; and seminar participants at AEA, Washington (Olin and Economics), U.S. Treasury (OFR), Richmond Fed, St. Louis Fed, Utah (Eccles), Johns Hopkins (Economics), Fed Board, Esade Workshop, Stanford Institute for Theoretical Economics (Causes and Consequences of Misallocation), Advances in Field Experiments (Chicago), Northeastern (Economics), George Mason (Economics), Boston (Questrom), Illinois (Gies), MIT (Sloan), European Winter Finance Conference (Megeve), George Washington (Economics), Barcelona Summer Forum (Financial Intermediation and Risk) for valuable suggestions. Washington University IRB (HRPO #202209011, *Capital Reallocation*).

When a bank earns higher returns at some branches than others, its headquarters can—in principle—reallocate capital until marginal returns are equalized. Whether internal capital markets actually achieve this is a fundamental question in corporate finance. Bank internal capital markets intermediate tens of trillions of dollars in the U.S. alone,[1] yet empirical tests of their efficiency have been hampered by a lack of internal data and, crucially, exogenous variation to measure how profits respond when capital is shifted across branches.

In this paper, we conduct a large-scale field experiment to test the allocative efficiency of a bank's internal capital market.[2] We study a large European bank whose approximately 500 branches share a common cost of funds, and use a randomized pricing experiment on approximately 50,000 loan applications to identify the marginal return on capital for each branch group—the change in profit per dollar of capital allocated. We cannot reject that the total lending scale is optimal, but we reject equalized marginal returns across branch groups, implying potential reallocation gains of 8.4% of profits across physical branches. We identify the bank's branch-uniform pricing policy—charging the same spread across all branches conditional on credit score—as an important impediment to price-based reallocation.

Our setting—capital allocation *within* a single organization—is distinct from the across-firm misallocation studied in the macroeconomics literature (e.g., Hsieh and Klenow (2009), Carrillo et al. (2023)). Within a firm, headquarters controls several allocation levers; in our setting, we experimentally vary one of them—loan pricing—and observe how capital deployment and profits respond. This makes allocative inefficiency more directly testable, more surprising, and more actionable: the inefficiency is not a market failure but self-imposed, and the question is not whether a social planner could improve on decentralized markets, but whether the firm can improve on itself.

A notable feature of our analysis is that *average* returns are misleading about the scope for reallocation. Using historical data, we document large and persistent dispersion in average returns across branches: the average return at top-decile branches is nearly four times that at bottom-decile branches. This raises concern about misallocation—too much capital allocated to low-return branches, and vice versa—but average returns do not reveal the marginal returns as capital is reallocated, and hence do not

[1] U.S. commercial banks alone hold over $23 trillion in assets; see Board of Governors of the Federal Reserve System (2026), Federal Reserve, H.8 release, *Assets and Liabilities of Commercial Banks in the United States*.

[2] To our knowledge, no prior field experiment has been designed to test allocative efficiency within a single firm's internal capital market.

by themselves measure feasible gains. The reason is that reallocating capital is costly: expanding lending in a branch requires lowering its price, which reduces the return on every loan that would have been issued at the original price—an inframarginal effect. Our experiment allows us to measure this cost directly.

In Section 2, we formalize this logic.[3] We define the marginal return as:

$$\mu = \frac{\Pi'(R)}{\Theta'(R)},$$

where $\Pi'(R)$ and $\Theta'(R)$ are the responses of profit and capital allocated to a change in price R. An *allocative inefficiency* exists if marginal returns are not equalized across branches: the bank could adjust prices to induce a reallocation and increase profits. For an unconstrained bank, optimality requires $\mu = 0$ for each branch. For a bank operating under a total capital constraint, allocative efficiency requires that μ equal a common value across branches (not necessarily zero), since the opportunity cost of capital deployed in one branch is a forgone loan in another. Using a second-order approximation, we link the cross-sectional dispersion of marginal returns to potential reallocation gains via a closed-form expression.

In Section 3, we describe the experiment. The bank employed risk-based pricing: the spread depended solely on the public credit score, and was therefore uniform across branches. The experiment randomly assigned applicants to receive a price either lower or higher than the baseline. Take-up responds strongly to price—the experimental price cut of $1.32pp$ APR raises take-up by $13.2pp$ (38%)—so the experiment randomly diverts capital toward or away from a branch: prices move quantities, and the quantity here is capital. Because the experiment includes both price cuts and hikes, it also allows us to estimate the curvature of the profit function without imposing a common functional form across branches.

In Section 4, we present the findings.[4] We estimate the weighted mean of marginal returns across physical (excluding online) branch groups to be 0.17% (s.e. 0.30%); hence, we cannot reject that the total lending scale is optimal. However, we reject equalized marginal returns across branch groups (p=0.002). Approximately 40% of branch groups (by lending volume) exhibit a positive marginal return, suggesting that capital should

[3]The approach circumvents the need for restrictive assumptions about demand and cost—and is agnostic: it does not require us to take a stance on the extent of optimizing behavior or the root cause of inefficiencies.

[4]We group the bank's approximately 500 branches into 21 branch groups (20 groups of physical branches plus the online branch) based on persistent unexpected overperformance to ensure sufficient statistical power within each group. All references to "branch groups" refer to these groupings.

increase, while the remainder exhibit negative marginal returns. The implied gross reallocation gain is 8.4% for physical branches and 10.7% of profits for all branches, including online. These gains arise predominantly from the dispersion of marginal returns: reallocation alone—holding total lending fixed—captures 95% of the unconstrained gain, highlighting the value of *reallocating* existing capital rather than expanding total lending. Our results are robust to accounting for cross-sell and origination costs, and winsorization.

It may be surprising that capital allocation is inefficient *within* an organization that, by virtue of its control rights, can reallocate capital at will in ways that external markets cannot. Moreover, the classic friction that the literature invokes to explain misallocation in internal capital markets—capital constraints at the division level, e.g., Stein (1997)—is absent here. This makes the failure to equalize marginal returns more puzzling.

We identify the bank's branch-uniform pricing policy as an important constraint on price-based reallocation: by charging the same spread across all branches, the bank forgoes the ability to use prices to equalize marginal returns. In our conversations with the bank, two concerns were frequently raised against branch-level pricing: strategic responses from customers manipulating branch choice, and perceptions of unfairness. Our reallocation gains provide a benchmark against which the costs of accommodating these concerns can be assessed. Notably, after reviewing our results, the bank began revising its practices to incorporate branch-level factors into the pricing decision.

Literature. Our paper contributes to the literature on the efficiency of internal capital markets in finance (see Lamont (1997); Shin and Stulz (1998); Gertner et al. (1994); Stein (1997); Scharfstein and Stein (2000); Whited (2001); Duchin and Sosyura (2013)).[5] Prior studies have used naturally occurring data to test for efficient allocation within conglomerates and non-financial firms; the banking literature documents that bank internal capital markets *exist* (see Peek and Rosengren (1997); Houston et al. (1997); Campello (2002); Cremers et al. (2011); Gilje et al. (2016); Cortés and Strahan (2017); Drechsler et al. (2017)), but evidence on whether they allocate capital *efficiently* has been limited.[6] We

[5] Stein (2003) surveys the identification difficulties—marginal investment opportunities are unobservable and q-proxies fail—that our experimental wedges are designed to overcome. Maksimovic and Phillips (2002) study how conglomerates allocate resources across divisions over the business cycle. Matvos and Seru (2014) estimate a structural model of internal capital markets. Giroud and Mueller (2015) use new airline routes as exogenous shocks and document reallocation within firms, but only if the firm is financially constrained. Giroud et al. (2022) estimate that within-firm distortions account for roughly 30% of the potential gains from equalizing marginal products across U.S. plants—about half among multi-plant firms.

[6] Indeed, within-bank empirical designs commonly invoke a frictionless internal capital market—funds flow across branches until marginal returns to lending are equalized—as a simplifying assumption (Drechsler et al., 2017). Our paper also relates to financial integration more broadly, e.g., Jayaratne and Strahan

contribute by conducting a field experiment to facilitate capital reallocation within a bank, directly measuring the marginal returns this literature can only proxy, testing whether marginal returns are equalized, and quantifying the gains.

Our paper also relates to the macroeconomics literature on misallocation of capital across firms (e.g., Hsieh and Klenow (2009), Gopinath et al. (2017), Whited and Zhao (2021), David et al. (2022), Bau and Matray (2023), Sraer and Thesmar (2023)). Two recent papers use experimental variation to measure misallocation: Carrillo et al. (2023) use construction procurement lotteries in Ecuador, and Hughes and Majerovitz (2024) use randomized grants to microenterprises in Sri Lanka. Both study misallocation *across* firms. Our contribution is to adapt these tools to study misallocation *within* a firm, where the planner's own allocation lever—loan pricing—is the instrument, and where the inframarginal cost of reallocating through price is first-order.

Finally, we contribute to the literature on uniform pricing—see (Nakamura, 2008; Cavallo et al., 2014; Hitsch et al., 2021; DellaVigna and Gentzkow, 2019; Daruich and Kozlowski, 2023) for retail and (Park and Pennacchi, 2009; Granja and Paixao, 2026; Dlugosz et al., 2024) for banks—as well as studies featuring randomized pricing experiments.[7] Uniform pricing is pervasive across industries and countries; our reallocation gains quantify the cost of this practice within a bank's internal capital market.

1 Environment and Data

1.1 Institutional Details

Collaborator. We collaborate with a bank headquartered and regulated in the E.U. Our analysis uses data exclusively from the bank's subsidiary in Türkiye. This subsidiary ranks in the top-10 by asset size and number of branches in the local market, with a customer base broadly representative of banked households. It offers a wide range of products, including unsecured loans, mortgages, checking and savings accounts, overdrafts, retirement accounts, brokerage, and insurance.

Unsecured loans. We focus on a single product—unsecured installment loans with

(1996); Landier et al. (2017).

[7] For example, DeFusco et al. (2022) and Aydin (2024). DellaVigna and Gentzkow (2019) show that most U.S. food, drugstore, and mass-merchandise chains price uniformly across stores. Granja and Paixao (2026) and Begenau and Stafford (2023) provide evidence that U.S. banks predominantly use uniform deposit pricing. Hurst et al. (2016) document price uniformity for mortgages securitized by GSEs, but not in the private market. Hazell et al. (2025) find that firms set the same nominal wage across establishments regardless of local marginal revenue products and calculate a model-implied profit cost of uniform wage setting of 2.5% for the median establishment.

fixed rates, fixed nominal payments in local currency, and maturities of up to 72 months. These loans account for two-thirds of total non-mortgage household debt outstanding, so the allocation we study is first-order for the lender. Borrowers commonly use these loans for debt consolidation, home improvements, health expenses, travel, education, or large purchases.

Table 1: Environment: Underwriting

Policy rule	Tier:	Prime	Mid	Sub
Application rejected if larger than				
Debt (Total)		100,000	75,000	50,000
Debt-to-income		10	8	6
Payment-to-income		50%	40%	30%

Note. Table displays, for illustration purposes, a hypothetical underwriting rule. The actual underwriting rule has many more policy rules.

Underwriting. At the time of application, applicants provide basic demographic information (e.g., education, employment, home ownership, income). They then declare a purpose for borrowing and request a loan size and maturity. Applicants are tiered based on scores. These tiers determine total debt, debt-to-income, and payment-to-income limits. For illustration purposes, Table 1 outlines a hypothetical underwriting policy. In this example, an applicant in the Prime tier is rejected if total debt is higher than 100,000 or debt-to-income is higher than 10 or payment-to-income is higher than 50%. Underwriting is segmented such that salary customers and retirees (with verifiable incomes) are subject to separate policies. Underwriting allows for managerial discretion, as branch managers can approve marginal applicants (i.e., those near the thresholds in Table 1), but they can *not* modify the price, size, or maturity.

Pricing. The loan price is determined by adding a spread to the cost of funds. The cost of funds depends on the loan type and currency, but not on the credit score. The spread depends solely on the credit score (via estimated cost-of-risk). Hence, pricing is *risk-based.* Table 2 shows the bank's rate sheet during the time of the experiment. All applicants, excluding a small fraction ($\approx$ 10%) who receive a *preferred* price (e.g., bank employees, salary customers, retirees), are assigned one of five spreads, as displayed in the Table. Table 3 shows that these five credit score bins in Table 2 alone explain more than 98% of the variation in the spreads. As the credit score is the sole input (and branch

is not an input), prices are uniform across branches.

Table 2: Environment: Bank's *Rate Sheet*

Score	**Cost of Funds** % APR	**Spread** % APR	**Rate** % APR
1600 to 1900	15.7	2.5	18.2
1400 to 1600	15.7	2.9	18.6
1250 to 1400	15.7	3.5	19.2
1150 to 1250	15.7	3.8	19.5
900 to 1150	15.7	4.2	19.9

Note. Table displays the actual rate sheet used by the bank during the experiment.

Table 3: Environment: Risk-based Pricing

	Spread		**Cost of Funds**	
Score—5 Bins	✓		✓	
Branch		✓		✓
R^2	0.983	0.018	0.005	0.023

Note. Table displays R^2 from a regression where the dependent variable is either the spread, or the cost of funds. 5 Credit Score bins are as in Table 2.

Internal Capital Market. Our setting is unambiguously an internal capital market in the classic control-rights sense (e.g., Alchian, 1969; Williamson, 1975; Gertner et al., 1994). The bank comprises approximately 500 branches. These branches face distinct applicant pools and local competition, differ in informational access, and retain some underwriting discretion at the margin. Most notably, the branches access a common pool of capital at a common cost of funds; the opportunity cost of the capital deployed in one branch is a forgone loan in another branch. Loans are financed on the bank's balance sheet; no loan is individually funded (or securitized) externally. There are no lending or deposit constraints at the branch level, and the movement of capital across branches is unrestricted. The headquarters holds authority over every allocation lever—prices,

underwriting, targets—and branch managers cannot modify the price.[8][9]

1.2 Returns to Capital—Measurement and Dispersion

Here, we introduce our data and returns to capital measures. We then use historical data to document persistent dispersion in average returns across branches.

1.2.1 Measurement

Data. We analyze historical data on all unsecured personal installment loans issued by the bank from 2015 to 2020. We construct the dataset by merging information on the applicant, the application, and loan performance. Applicant data includes credit score and income. Loan characteristics include the price, maturity, size (face value), and cost of funds. Performance data includes information on the timing of default (i.e., the date the loan became non-performing) and recovery after default (within the first 12 months).

Loan-level Returns. We measure the loan-level return as a net present value, normalized by face value. This gives a profit per dollar lent. We denote this return as π. We calculate net present value by discounting payments at the loan-specific cost of funds c, accounting for missed payments via the timing of default and subsequent recovery, both recorded as realized in administrative data without error.[10]

Branch-level Returns. We calculate branch-level returns by averaging loan-level returns at the branch level. First, we calculate branch-year returns. We then average these returns at the branch level to calculate long-run (2015-2020) branch returns.

1.2.2 Dispersion

We document that branch-level average returns exhibit substantial and persistent dispersion.

Focusing on 2015-2020, the branch-level return range (10th to 90th percentile) is

[8] Capital is (softly) budgeted at the start of the fiscal year through volume targets set at the branch level, as well as across segments and products. Branch managers are incentivized and solicit applications to achieve these targets. Otherwise, there are no explicit costs associated with allocating too much or too little capital to any particular branch. Because capital is not constrained at the branch level, our setting does not speak to winner-picking theories of allocation under scarcity, as in Stein (1997).

[9] Unlike other jurisdictions (e.g., India, Pakistan, and the Philippines), where regulators can meaningfully constrain banks' physical footprints—through branch authorization rules and, in some cases, approval requirements for closures or location changes—Türkiye does not impose comparable constraints on branch siting or closure.

[10] None of our average or marginal return calculations contain an accounting-based ingredient. Hence, dispersion in measured returns cannot reflect measurement error in loan outcomes (e.g., Bils et al., 2021). Further details are provided in Appendix A.

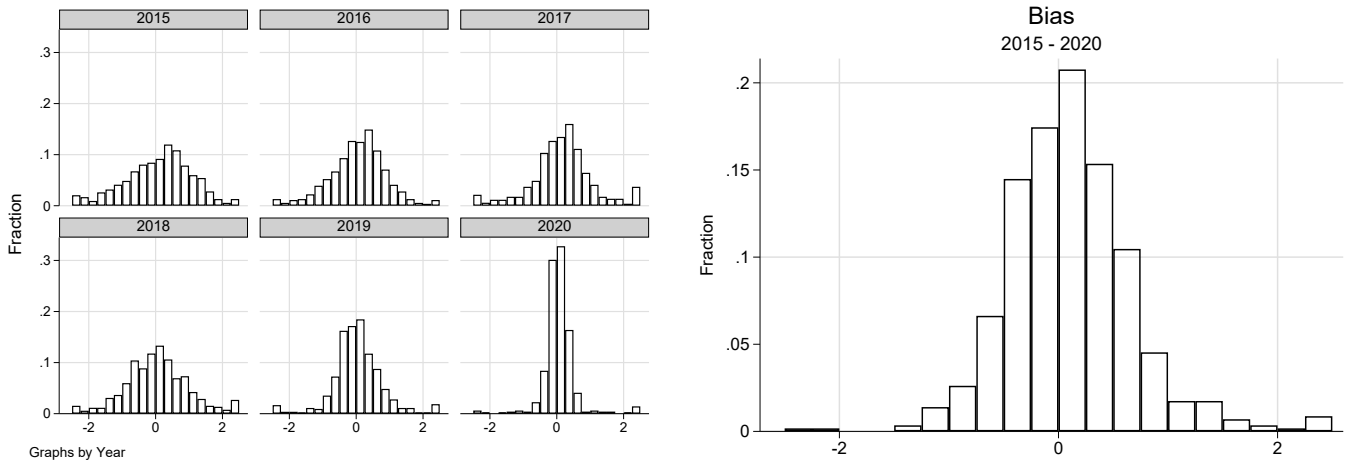
1.45pp for realized returns. The difference in realized returns between any two randomly chosen branches is 0.73pp—about half the size of the average return.

Naturally, differences in branch-level returns may be due to observables (e.g., credit score, loan size, and maturity). To isolate unexplained branch-level returns, we extract branch-year effects f_{bt} from a regression that purges loan-level realized returns π_j^R of observables X_j:[11]

$$\pi_j^R = g(X_j) \cdot f_t + f_{bt} + \varepsilon_j \tag{1}$$

Figure 1 displays the histogram of these branch-year effects f_{bt} by year. The branch effect explains 37% of the variation in these branch-year effects across six years. A branch in the bottom quartile has a 49% chance of remaining in the bottom quartile after one year, and a 40% chance of remaining there after three years. These patterns suggest that the branch effect on realized returns is meaningfully persistent. Indeed, these branch-year effects also predict average returns out of sample—on loans issued during the experimental period, with returns realized up to three years later.

Figure 1: Historic Data: Average Returns—Branch Overperformance



Note. The left figure displays the histogram of branch-year effects f_{bt} for 2015 to 2020 from Equation (1). The units are % of face value. The right figure displays the histogram of $Overperform_b = \sum_t f_{bt}$.

We then define $Overperform_b = \sum_t f_{bt}$—*persistent unexpected overperformance*—by aggregating these branch-year effects over time. The right panel of Figure 1 displays the

[11] X consists of credit score, loan size, and maturity. The score is developed by the public credit bureau (*Kredi Kayıt Bürosu*) in collaboration with Experian. Specifically, g interacts credit score (8 bins, plus one for null scores), size, and maturity (4 bins each), yielding $9 \cdot 4^2 = 144$ bins, and includes credit score as an additional regressor. We also interact $g(X_j)$ with calendar year effects f_t to allow the impact of observables to vary over time.

histogram of $Overperform$ across branches. Because *Overperform* is estimated entirely from pre-experiment data, the branch grouping we use in Section 4 is predetermined with respect to the experiment. $Overperform$ gives us a natural sorting variable for our reallocation exercise in Section 4.

A natural concern is that overperformance is compensation for systemic risk—a mean-variance bank equalizes risk-adjusted returns.[12] Table 4 shows the opposite—overperformance is *negatively* correlated with cyclicality and risk adjustment amplifies the dispersion.[13]

Table 4: Overperformance and Cyclicality

Branch	β	p
Cyclicality	-0.30	<0.001
Growth $_{\log}$	0.045	0.17

Note. $Cyclicality_b$ is the slope from regressing branch-year returns on the contemporaneous average return across all branches, estimated over 2015-2020. Growth is the percent change in lending over the year.

Overall, the dispersion in returns *invites the interpretation* that the bank may be allocating too much capital to low-return branches, and vice versa. However, *average* returns do not reveal much about the *marginal* (next-best) returns for each branch. Marginal returns also depend on demand and on how demand and costs change as capital is reallocated. Next, we conceptually characterize these marginal returns.

2 Conceptual Framework

This section develops our approach: we identify marginal returns using supply variation and link the estimates to tests of allocative efficiency and to potential reallocation gains. Our framework applies broadly to settings where capital is allocated among competing segments (e.g., branches, divisions, products) via different levers (e.g., prices, cost of funds, underwriting thresholds), see Table 15. Here, we focus on reallocation across branches using pricing variation.

[12] For a diversified bank, shareholders hold the market, so only systematic risk is priced; hence low-return branches must be low-beta. For an undiversified or capital-constrained bank à la Froot and Stein (1998), headquarters is mean-variance over its own book; hence low-return branches must have low covariance with the bank's book. Either way, low-return branches should be less cyclical.

[13] A related concern is that what appears as underperformance is an investment in growth—if so, underperforming branches should be growing faster. We find overperformance is uncorrelated with branch growth, and if anything, overperforming branches grow faster.

Consider a bank lending across $b \in \{1, \ldots, B\}$ branches. These branches share an internal capital market and a common cost of funds, c. Demand and cost of risk are heterogeneous across the branches. Otherwise, there is no heterogeneity in observables (e.g., size, maturity, credit score).

Let R_b denote the price—the interest rate on the loan—and $q_b(R_b)$ the average cost of risk conditional on take-up (missed payments per dollar lent). We assume q is increasing in R—reflecting adverse selection and moral hazard (which we do not distinguish).

As in Section 1, the average return (i.e., *average profit per capital allocated*) is:[14]

$$\pi_b(R_b) = R_b - c - q_b(R_b). \tag{Average}$$

π is defined only for issued loans and does not capture the bank's ability to convert applications into loans. A bank may accept a lower profit from issued loans if increased demand compensates for the reduction in R. Let $\Theta_b(R_b)$ represent demand (take-up). We assume Θ is decreasing in R.

Faced with a loan application, the bank's objective is to maximize concave profits—the product of demand and average return:

$$\begin{aligned}\Pi_b(R_b) &= \Theta_b(R_b)\ \pi_b(R_b) \\ &= \Theta_b(R_b)\ \Big(R_b - c - q_b(R_b)\Big).\end{aligned}$$

For expositional simplicity, we first drop the subscript b, and study an *unconstrained* bank—i.e., no total capital constraint—allocating capital to a single segment.

The optimal price, denoted R^*, is chosen as:[15]

$$R^* = \arg\max_R \Theta(R)\Big(R - c - q(R)\Big),$$

[14] Later, we generalize π to incorporate *cross-sell* o and cost of origination ζ.

[15] Although our approach and the experiment operate through prices, both are inherently tied to the capital allocation problem: because the demand curve is strictly downward-sloping, deploying more capital requires lowering the price. Indeed, holding the technology (e.g., underwriting, advertising) constant, the marginal return μ we identify from price variation is equivalent to the marginal return of the capital allocation problem against the inverse demand. See early work by Weitzman (1974) and Table 15 for a discussion of this equivalence.

with the first-order condition:

$$\Pi'(R^*) = \Theta'(R^*)\left(R^* - c - q(R^*)\right) + \Theta(R^*)\left(1 - q'(R^*)\right) = 0. \qquad \text{(FOC)}$$

In the first-order condition, the first term represents the extensive margin effect of a change in price. This is the main allocation lever. Each issued loan earns a positive average return π. Because the demand curve is downward-sloping, allocating more capital requires lowering prices. Lowering prices increases demand, where each new loan brings an additional average return (and vice versa for raising prices). The second term represents the intensive margin—the *inframarginal* effect—the cost of allocating more capital: lowering the rate also affects the returns on loans already taken up, through changes in price and the cost of risk. The inframarginal effect, not a shortage of positive-NPV loans, is what keeps the bank from lending more. Optimality is achieved by balancing these two effects.

2.1 Marginal Returns

Define a *wedge* μ such that the observed price $\hat{R}$ maximizes a *distorted* objective:

$$\hat{R} = \arg\max_R \Theta(R)\Big(R - c - q(R) - \mu\Big),$$

and satisfies the distorted first-order condition:

$$\Theta'(\hat{R})\left(\hat{R} - c - q(\hat{R}) - \mu\right) + \Theta(\hat{R})\left(1 - q'(\hat{R})\right) = 0. \qquad \text{(FOC-}\mu\text{)}$$

μ rationalizes the observed price $\hat{R}$ and may be interpreted as a *misperception* about cost of funds c.[16]

Notice that (FOC) evaluated at $\hat{R}$ is:

$$\Pi'(\hat{R}) = \mu\,\Theta'(\hat{R}).$$

The change in profit at $\hat{R}$ equals the product of μ and the slope of demand. Hence, μ can

[16] It could also be interpreted as a misperception about cost of risk q or average return π. For example, if μ=1%, then the observed price $\hat{R}$ could be rationalized if (or chosen under the misperception that) the cost of funds c or cost of risk q is $1pp$ higher, or the average return π is $1pp$ lower.

be estimated, using price variation around $\hat{R}$, as:

$$\begin{aligned}\mu &= \frac{\Pi'(\hat{R})}{\Theta'(\hat{R})} \\ &= \pi(\hat{R}) + \frac{\Theta(\hat{R})}{\Theta'(\hat{R})}\left(1 - q'(\hat{R})\right). \qquad \text{(Marginal)}\end{aligned}$$

μ is the *change in profit per marginal capital allocated*—i.e., the *marginal return*. The marginal return depends on the average return, π, as each new loan brings additional profit. The marginal return also depends on the costs of reallocating, through the level of demand Θ and how demand and costs change with prices—Θ' and q'.

The marginal return is high if the average return is high and the cost of reallocation is low; the cost of reallocation is low if a small price change shifts a lot of capital while repricing few existing loans, and if the price change lowers cost of risk.[17]

In general, the marginal return differs from the average return if and only if the average return moves with scale. The average return may move with scale for three reasons: the price obtainable at a given quantity (market power), the composition of who is served (selection), or the unit cost of the technology (returns to scale).[18] In our setting, the first two channels are explicit: deploying more capital requires lowering the interest rate against a downward-sloping demand, which shrinks the margin on every loan—including those the bank would have made anyway (market power)—and the cost of risk $q(R)$ moves with the price as the composition of the borrower pool changes (selection).

We now define our optimality conditions, summarized in Table 5.

[17] Two limits are instructive. As $\Theta' \to -\infty$ (perfect competition), $\mu = \pi$: reallocation is costly only because the bank faces a downward-sloping demand. If the cost-of-risk slope is sufficiently high, at $q' = 1$, the improvement in the pool exactly pays for the price cut, and again $\mu = \pi$.

[18] Formally,

$$\mu = \Pi'(R)/\Theta'(R) = \pi + \frac{\Theta}{\Theta'}\pi' = \pi + \Theta\frac{d\pi}{d\Theta}.$$

In the misallocation literature, average revenue products are proportional to marginal revenue products under CES demand and common technology (e.g., Hsieh and Klenow, 2009). In selection markets, marginal cost equals average cost absent selection (no unobservable risk correlated with demand); the difference is composition (e.g., Einav and Finkelstein, 2011). In q-theory, marginal q equals average q under constant returns (Hayashi, 1982); the difference is the curvature of the production technology—and, empirically, measurement error further severs average q from marginal q (Erickson and Whited, 2000).

2.2 Unconstrained Optimality

Unconstrained optimality (i.e., scale optimality, with no total capital constraint) requires that the chosen price satisfy the first-order condition, corresponding to the null hypothesis $H_0 : \mu = 0$. When $\mu \neq 0$, an unconstrained bank can increase profits by adjusting prices.[19] A positive μ indicates that capital allocation should increase. Because the demand curve is downward-sloping, this requires lowering the price. Conversely, a negative μ implies that profits would increase if capital allocation decreased (and price increased).

Table 5: Conceptual Framework: Optimality Conditions

Unconstrained Optimality i.e., *Scale Optimality*	$\mu = 0$
Constrained Optimality i.e., *Allocative Efficiency*	$\text{var}(\mu_b) = 0$
Reallocation Gain	$\mathcal{G}^C > 0$

2.3 Constrained Optimality

Now consider the case where the bank faces a total capital constraint, such that:

$$\sum_b \Theta_b(R_b) \leq \bar{\Theta}.$$

The prices are then chosen as:

$$\begin{aligned} \{R_1^*, \dots R_B^*\} &= \arg \max_{R_1,\dots,R_B} \quad \sum_b \Theta_b(R_b)\Big(R_b - c - q_b(R_b)\Big) \quad \text{s.t.} \quad \sum_b \Theta_b(R_b) \leq \bar{\Theta} \\ &= \arg \max_{R_1,\dots,R_B} \quad \sum_b \Theta_b(R_b)\Big(R_b - c - q_b(R_b) - \mu\Big). \end{aligned}$$

Notice that the presence of a total capital constraint gives rise to the same *distorted* first-order condition as (FOC-μ). Now, μ is the Lagrange multiplier on the total capital

[19] The same adjustment can equivalently run through the cost of funds c rather than the spread $R - c$.

constraint.[20] When the constraint binds, the effective cost of capital becomes $c + \mu$ for all uses. While μ need not equal zero, it should be common across uses. Any *dispersion* in the marginal return indicates allocative inefficiency across uses.

2.4 Reallocation Gains

To quantify potential gains from reallocating capital across uses (i.e., holding total allocation constant), we derive an analytical formula. The bank's reallocation gain due to the *dispersion* in marginal returns, in the presence of the constraint, is:

$$\begin{aligned}\mathcal{G}^C &= \Pi(R^*) - \Pi(\hat{R}) \\ &\approx -\sum_b \frac{1}{2}\,(\mu_b - \bar{\mu})^2\,\xi_b \quad \text{where } \bar{\mu} = \frac{\sum_b \xi_b\,\mu_b}{\sum_b \xi_b} \text{ and } \xi_b = \frac{\Theta_b'(\hat{R}_b)^2}{\Pi_b''(\hat{R}_b)}. \end{aligned} \tag{Gain}$$

We present the proof in Section 2.6, which follows from a second-order Taylor expansion of Π around $\hat{R}$. The gain is the triangle area implied by the local approximation: it equals one-half the reallocation $(\Theta_b^* - \hat{\Theta}_b)$ times the excess marginal-return wedge $(\mu_b - \bar{\mu})$, summed across uses. $\bar{\mu}$ is a weighted mean of the marginal returns, with weights chosen to ensure that the total capital constraint is satisfied.

This expression summarizes, in terms of marginal returns, the gains from reallocating capital across uses while holding total allocation constant. This is done by adjusting prices such that μ were common across uses at the level $\bar{\mu}$.

The reallocation gain depends on the level of the marginal return μ, the slope of demand Θ', and the curvature of profits Π'' (all evaluated at $\hat{R}$). If μ_b were common across uses at the level $\bar{\mu}$ there would be no gain. Otherwise, reallocation gains are quadratic in μ—small distortions around the optimum result in negligible gain, but gains grow rapidly with μ.[21]

[20] The Lagrangian associated with the constrained optimization problem is:

$$\mathscr{L} = \sum_b \Theta_b(R_b)\Big(R_b - c - q_b(R_b)\Big) + \mu\left[\bar{\Theta} - \sum_b \Theta_b(R_b)\right]$$

[21] Under linear demand and cost curves (i.e., $\Theta'' = q'' = 0$), the gain simplifies to $\frac{1}{4}\mu^2\frac{-\Theta'}{1-q'}$. The loss is monotonically decreasing in the curvature of profits — a given dispersion in marginal returns is worth less the faster marginal profit deteriorates as prices move. Hence, dispersion in marginal returns is necessary for a gain but not sufficient: the realizable gain is also governed by the curvature of profits—as $\Pi'' \to -\infty$, dispersion persists but the gain vanishes.

Table 6: Conceptual Framework: Gain and Reallocation

μ	Θ'	Π''	$R^* - \hat{R}$		$\Theta^* - \hat{\Theta}$		$\Pi^* - \hat{\Pi}$
% of FV			pp	% (of $\hat{R}$)	pp	% (of $\hat{\Theta}$)	% of $\hat{\Pi}$
0.5%	-10	-10	-0.5	-2.5	5	20	2.5%
1%	-10	-10	-1	-5	10	40	10%
2%	-10	-10	-2	-10	20	80	40%
1%	-5	-10	-0.5	-2.5	2.5	10	2.5%
1%	-10	-20	-0.5	-2.5	5	20	5%
-1%	-10	-10	1	5	-10	-40	10%
1%	-20	-10	-2	-10	40	160	40%

Note. Table displays the implied price adjustment, reallocation, and gain as a function of the underlying moments, under the assumption of $\hat{\Pi} = 0.5\%$, $\hat{R} = 20\%$ APR, $\hat{\Theta} = 25\%$. For example, a marginal return $\mu = 1\%$, with $\Theta' = -10$ and $\Pi'' = -10$, implies a price change of -1pp, a reallocation of 10pp, and a gain of 10% of $\hat{\Pi}$.

2.5 Taking Stock

Allocative efficiency requires that marginal returns be equalized across uses, though the common value need not be zero. When marginal returns are not equalized, the bank can raise profits—given its constraints and information—by adjusting prices to reallocate capital, holding the total fixed, until the dispersion is eliminated.

Table 7 summarizes the concepts used throughout the paper—average return, marginal return, reallocation, and gain. Note that for two uses with the same demand level, demand slope, and cost-of-risk slope, any difference in marginal returns is entirely due to the difference in average returns.

2.6 Proofs

We start by deriving an analytical formula for the reallocation gains of an *unconstrained* bank, in two steps.

First, we derive the relationship between μ and the deviation of the observed price $\hat{R}$ from the optimal level R^*:

$$R^* - \hat{R} \approx -\mu \frac{\Theta'(\hat{R})}{\Pi''(\hat{R})}, \qquad (R^* - \hat{R})$$

Table 7: Conceptual Framework: Average, Marginal, Reallocation, Gain

	Symbol	Formula	Definition
(Average)	π_b	$= \frac{\Pi_b}{\Theta_b} = R_b - c - q_b$	*Average profit per capital allocated*
(Marginal)	μ_b	$= \frac{\Pi'_b}{\Theta'_b} = \pi_b + \frac{\Theta_b}{\Theta'_b}(1 - q'_b)$	*Change in profit per marginal capital allocated*
(Reallocation)	$\Theta^*_b - \hat{\Theta}_b$	$= -(\mu_b - \bar{\mu})\, \frac{\Theta'^2_b}{\Pi''_b}$	*Optimal minus observed capital allocation*
(Gain)	$\Pi^* - \hat{\Pi}$	$= -\sum_b \frac{1}{2}\, (\mu_b - \bar{\mu})^2\, \frac{\Theta'^2_b}{\Pi''_b}$	*Optimal minus observed profit*

which follows from a first-order Taylor expansion of Π' around $\hat{R}$:

$$\Pi'(R^*) = \Pi'(\hat{R}) + \Pi''(\hat{R})\ (R^* - \hat{R}) + O((R^* - \hat{R})^2),$$

where $\Pi'(R^*) = 0$ and $\Pi'(\hat{R}) = \mu\ \Theta'(\hat{R})$.

Second, the expression for the unconstrained reallocation gains follows from the second-order Taylor expansion of Π around $\hat{R}$:

$$\Pi(R^*) = \Pi(\hat{R}) + \Pi'(\hat{R})\ (R^* - \hat{R}) + \frac{1}{2}\Pi''(\hat{R})\ (R^* - \hat{R})^2 + O((R^* - \hat{R})^3),$$

where $\Pi'(\hat{R}) = \mu\ \Theta'(\hat{R})$ and $R^* - \hat{R} \approx -\mu\ \frac{\Theta'(\hat{R})}{\Pi''(\hat{R})}$.

Hence, the unconstrained reallocation gain is given by:

$$\mathcal{G}^U(\mu) = -\frac{1}{2}\ \mu^2\ \frac{\Theta'(\hat{R})^2}{\Pi''(\hat{R})}.$$

The total reallocation gain is the sum across uses of the gain $\mathcal{G}_b$:

$$\mathcal{G}^U_{\text{Total}} = \sum_b \mathcal{G}^U_b(\mu_b).$$

The analytical formula for the reallocation gains of a *constrained bank* follows anal-

ogously. It only remains to be shown that the total capital constraint is satisfied.

First, note that the difference between the observed allocation $\hat{\Theta}$ and the optimal level Θ^* is, to the first order:

$$\Theta^* - \Theta(\hat{R}) \approx \Theta'(\hat{R})\left(R^* - \hat{R}\right) \approx -\mu \frac{\Theta'(\hat{R})^2}{\Pi''(\hat{R})}.$$

In constrained reallocation, we replace μ_b with $\mu_b - \bar{\mu}$:

$$\Theta_b^* - \Theta_b(\hat{R}_b) \approx -\left(\mu_b - \bar{\mu}\right) \frac{\Theta_b'(\hat{R}_b)^2}{\Pi_b''(\hat{R}_b)}. \qquad \text{(Reallocation)}$$

The total change in capital allocation is thus $\sum_b (\bar{\mu} - \mu_b) \frac{\Theta_b'(\hat{R}_b)^2}{\Pi_b''(\hat{R}_b)}$, which equals zero by the construction of $\bar{\mu}$, as required.

3 Experiment

We now describe the randomized experiment used to identify the marginal returns.

Conceptually, the goal is to induce more or less capital allocation by randomizing supply. Econometrically, this approach requires an instrument $\mathbb{Z}$ that is correlated with supply but uncorrelated with demand or costs. The experiment, which randomizes prices at a level either lower or higher than the baseline, does exactly that.

In this section, we present descriptive evidence on *average* effects. In Section 4, we estimate marginal returns at more granular levels.

Design. The experiment involves 50,042 applicants approved for underwriting.[22] Applicants were randomly assigned to three experimental groups based on the final digits of the customer identification number. These three groups were offered different prices. The baseline group (80%) received the risk-based price displayed in Table 2. The low-price (L) group, consisting of 10% of applicants, was offered a price reduction. The high-price (H) group, also consisting of 10% of applicants, was offered a price hike. Com-

[22] The experiment was conducted by the bank's treasury, with the stated goal of estimating the average slope of demand. The research team was not involved in the design or implementation. There was no explicit participation choice; therefore, applicants were not aware they were participants in a controlled trial. The experiment ran from November 2020 to January 2021. During the relevant period of the COVID-19 pandemic in 2020, Türkiye did not impose lockdowns during weekday business hours for the general population, opting instead for overnight and weekend curfews. Essential businesses, including banks, generally remained open during daytime hours. Nevertheless, a larger share of applications was finalized online.

Table 8: Experiment: Summary Statistics

Variable		unit	N	mean	s.d.	p10	p50	p90
Panel A: Applicant								
$\mathbb{Z}^L$			50,042	0.10	0.30	0	0	0
$\mathbb{Z}^H$			50,042	0.10	0.30	0	0	0
Metro			50,042	0.59	0.49	0	0	1
Score (Credit Bureau)			50,042	1,407	237	1,077	1,415	1,728
Income		TRY	42,639	6,125	10,188	2,300	4,032	10,000
Loans (Existing)		Count	50,042	1.6	1.7	0	1	4
Payments (Existing)		TRY	50,042	998	1,887	0	376	2,653
Debt (Existing)		TRY	50,042	24,504	41,560	0	10,000	66,595
Panel B: Loan								
Size	FV	TRY	50,042	29,591	27,879	6,500	24,000	50,000
Maturity	T	Months	50,042	29.0	9.4	12	36	36
Cost of Funds	c	% APR	50,042	15.67	0.7	14.9	16.2	16.2
Spread	$R - c$	% APR	50,042	3.17	0.9	2.3	2.9	4.2
Price	R	% APR	50,042	18.84	1.1	17.4	18.7	20.4
Panel C: Outcomes								
Demand	Θ	%	50,042	34.2	47.4	0	0	100
Cost of Risk	q	% of FV	17,096	1.72	9.84	0	0	0.12
Average Return	π	% of FV	17,096	1.73	9.81	1.29	3.53	5.40
Profit	Π	% of FV	50,042	0.59	5.79	0	0	4.08
Cross-sell	o	% of FV	17,096	0.25	1.37	0	0	0.56

pliance with the assigned price is perfect: the price is displayed on the application, and neither the branch nor the applicant could change it.

We denote the assignment to the low- and high-price groups as $\mathbb{Z}^L = 1$ and $\mathbb{Z}^H = 1$, respectively. Hence, the experimental procedure provides two randomized instruments, $\mathbb{Z}^L$ and $\mathbb{Z}^H$, for econometric evaluation. Importantly, this variation allows us to observe behavior in response to both price increases and decreases.[23]

Summary Statistics. Table 8 displays summary statistics. Panel A displays applicant characteristics. Panel B displays loan characteristics. The average price, cost of funds, and spread are 18.8%, 15.7%, and 3.2%. The average maturity is 29 months and the average requested face value is 29,591 TRY, about six months of income. The maximum maturity

[23] This is a key ingredient in our reallocation gains calculation as it enables the nonparametric estimation of the curvature of the profit function, Π''.

is three years, with two-thirds of all loans having a maturity of three years.

First, applicants decide whether to take up the loan (demand). A total of 17,096 applications (34.2%) were taken up and issued. Then, those who take up decide whether to miss payments (cost of risk). Among takers, the cost of risk is 1.72% of the face value. This yields an average realized return of 1.73% of the face value for issued loans and a profit of 0.59% of face value (1.73% $\times$ 34.2%). Our loan-level outcome variables are demand Θ and cost of risk q.

Table 9: Experiment: Covariate Balance

	Metro	**Score**	**Income** TRY	**Loans** Count	**Payments** TRY	**Debt** TRY	**Term** Months	**Cost of Funds** % APR
$\mathbb{Z}^L$ Low	-0.003	21.8	232	-0.01	-8.7	633	-0.3	0.003
	(0.007)	(3.6)	(162)	(0.03)	(28)	(624)	(0.14)	(0.01)
$\mathbb{Z}^H$ High	-0.012	-15.2	43	-0.02	-13.9	441	-0.2	0.04
	(0.007)	(3.6)	(168)	(0.03)	(29)	(630)	(0.14)	(0.01)
Baseline	0.59	1407	6,106	1.63	1,000	24,398	29.1	15.7
	(0.003)	(1.2)	(55)	(0.01)	(9)	(207)	(0.05)	(0.003)
N	50,042	50,042	42,639	50,042	50,042	50,042	50,042	50,042
p	0.27	$<$0.01	0.33	0.73	0.86	0.50	0.10	$<$0.01
$\%^L$	-0.5	1.6	3.8	-0.6	-0.9	2.6	-0.9	0.02
$\%^H$	-2.1	-1.1	-0.7	1.0	-1.4	1.3	-0.6	0.3

Balance. Table 9 displays the covariate balance table.

Identification. The randomized assignment to the treatment group, $\mathbb{Z}$, is independent of observable and unobservable determinants of demand and costs. Note that although the assignment to the treatment group, $\mathbb{Z}$, is random, the magnitude of the price increase is not randomized. In the analysis, we always restrict the variation used only to the random assignment of a participant, $\mathbb{Z}$.

Other Margins. Throughout the analysis, we consider loan terms (e.g., size, maturity) as fixed.[24] As price randomization is after underwriting approval, there is no scope for discretion on behalf of branch managers to offset the first-stage or affect other margins.

[24]Table 16 shows that a small fraction of applicants change the initial size or maturity declared at the application. For the few who do change, the change is unrelated to the price. We define an applicant as a changer in two ways: if the final contract size or term is different by more than 1% or 10%. According

Table 10: Experiment: Estimation Framework

	Z Instrument	X Regressor	Y Outcome	Specification
First-stage	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	R	$Y_i = \alpha + \theta^L \mathbb{Z}_i^L + \theta^H \mathbb{Z}_i^H + \epsilon_i$
Reduced Form	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	Θ or q	$Y_i = \alpha + \theta^L \mathbb{Z}_i^L + \theta^H \mathbb{Z}_i^H + \epsilon_i$
Treatment Effect	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	R	Θ or q	$Y_i = \alpha + \beta X_i + \epsilon_i$

Estimation Framework. Table 10 summarizes the estimation framework used to analyze the randomized controlled trial. First, we study the effects of the experiment on prices (first-stage), later on demand, cost of risk, and profit.

In the first stage, the left-hand-side variable is R, the price. In the reduced form specifications, the left-hand-side variable is either demand Θ or cost of risk q. The instruments $\mathbb{Z}$ are orthogonal to measurement error, omitted variables, and the error (e.g., factors that affect demand or costs for reasons other than the price).

To estimate a treatment effect within a two-stage least-squares (2SLS) framework, we replace the right-hand-side variable with the price R.[25]

Average Effect Estimates. Table 11 displays the first-stage and reduced form specifications and compares the average change in price and outcomes between the treatment group and the control group using ordinary least squares (OLS).

Focusing on the first-stage estimates in Table 11, the baseline group receives a price of 18.89% APR (corresponding to a spread of 3.22% APR), which is $1.32pp$ lower for group $\mathbb{Z}^L$ and $0.92pp$ higher for group $\mathbb{Z}^H$. The F-statistic for the first stage is 7,286. As the cost of funds is balanced across the groups, the difference in the price comes entirely from the spread.

The first column of reduced form estimates in Table 11 shows a downward sloping demand curve. The average demand for the baseline group is 34.5%, which is $13.2pp$

to these measures, 5.9% (1,022) to 8.1% (1,386) of the applicants change their term, and 4.8% (813) to 5.8% (976) change their loan size. However, neither the propensity to change term or size nor the magnitude of the change for changers varies across experimental groups.

[25] To estimate a treatment effect for price cuts, we exclude $\mathbb{Z}^H = 1$ and use $\mathbb{Z}^L$ as an instrument, and vice versa to estimate a treatment effect for price hikes. As the model is exactly identified, the treatment effect obtained using the linear two-stage least-squares (2SLS) framework corresponds to the Wald estimate, calculated as the ratio of the reduced form and first-stage effects.

Table 11: Experiment: First-stage and Reduced Form Estimates

	First-stage		Reduced Form	
	R **Price**	$R-c$ **Spread**	Θ **Demand**	q **Cost of Risk**
	% APR	% APR	%	% of FV
$\mathbb{Z}^L$ Low	-1.32	-1.32	13.2	-0.68
	(0.01)	(0.01)	(0.7)	(0.22)
Constant	18.89	3.22	34.5	1.81
	(0.005)	(0.004)	(0.2)	(0.08)
$\mathbb{Z}^H$ High	0.92	0.88	-16.6	0.41
	(0.01)	(0.01)	(0.7)	(0.35)
N	50,042	50,042	50,042	17,096
F	7,286	12,603		

(38%) higher for group $\mathbb{Z}^L$ and $16.6pp$ (48%) lower for group $\mathbb{Z}^H$. In terms of a treatment effect, a $1pp$ APR decrease in prices (5% off a base of 18.89%) increases demand by $\frac{13.2}{1.32}$=$10pp$ (29% off a base of 34.5%). Conversely, in response to a $1pp$ APR increase in prices, demand decreases by $\frac{16.6}{0.92}$=$18.1pp$ (52%). Demand is thus concave—it falls more steeply when prices rise than it rises when prices fall. Pooling the two arms (weighting by group shares), the linear model yields an average slope of -12.7. In elasticity terms, the demand slope implies an elasticity of $-\Theta'\frac{R}{\Theta} = 12.7 \times \frac{18.89}{34.5} \approx 7.0$ at the control price. This is the extensive margin of the reallocation exercise.

The second column of reduced form estimates in Table 11 shows an upward sloping cost of risk curve. The average cost of risk for the baseline group is 1.81% of FV, which is 0.68% of FV (38%) lower for group $\mathbb{Z}^L$ and 0.41% of FV (23%) higher for group $\mathbb{Z}^H$. Hence, in response to a $1pp$ APR decrease in prices, cost of risk decreases by $\frac{0.68}{1.32}$=0.52% of FV. Conversely, in response to a $1pp$ APR increase in prices, cost of risk increases by $\frac{0.41}{0.92}$=0.45% of FV. This yields a linear cost-of-risk, with an average slope of 0.47. In elasticity terms, the cost-of-risk slope implies an elasticity of $q'\frac{R}{q} = 0.47 \times \frac{18.89}{1.81} \approx 4.9$ at the control price. This is the cost-of-risk component of the intensive margin of the reallocation exercise.

Table 12: Experiment: Marginal Return Ingredients

μ	π	$\frac{\Theta}{\Theta'}(1-q')$	Θ	Θ'	q'
Marginal =	**Average +**	**Inframar.**	**Demand**	**Demand**	**Cost of Risk**
	Level		Level	Slope	Slope
0.28	1.71	-1.42	34.2%	-12.7	0.47

Note. The average return π and demand Θ are sample means Table 8. The slopes Θ' and q' are experimental estimates Table 11. The inframarginal term is calculated as in (Marginal).

4 Results

We now operationalize the framework developed in Section 2.

Table 12 displays the marginal return ingredients. The average return is 1.71% per dollar lent. Yet, allocating more capital (by lowering prices) yields an inframarginal effect of -1.42%—$= \frac{34.2\%}{-12.7}(1-0.47)$—where 34.2% is the average demand, and -12.7 and 0.47 are the average demand and cost slopes. This gives a marginal return of only 0.28% per dollar lent. Hence, the bank's *average* pricing and total allocation are close to optimal.[26]

Yet, these estimates mask dispersion. Next, we estimate heterogeneous marginal returns at more granular levels. Then, we calculate the cross-sectional dispersion of these marginal returns and quantify the reallocation gains.

Binning. The experiment involves 50,042 loan applications and 17,096 loans issued across approximately 500 branches. To estimate marginal returns and reallocate across them, we group loans into bins. Binning ensures a sufficient number of observations per bin and reduces statistical uncertainty in estimated demand and cost functions.

The choice of the number of bins B is a bandwidth problem. Too few bins understate the gain by lumping together branches with different μ. Too many bins leave few loans per bin, so bin-level estimates are noisy. We choose B=20, plus a separate bin for the online branch. This choice is at the conservative edge, yielding an average $\approx$ 2,500 applicants per bin: it concedes a partition loss of second order—which can only understate

[26] Incorporating the cost-of-risk response is essential to our measurement: with $q' = 0.47$, the improvement in the risk pool lowers the inframarginal effect by roughly half. A demand-only model, with $q' = 0$, would imply a marginal return of -0.98%—the wrong sign—and would counsel raising prices and shrinking the book precisely where capital should expand.

the gain—in exchange for well-estimated bin-level slopes.[27]

We bin branches based on $Overperform$—*persistent unexpected overperformance*, discussed in depth in Section 1.2 and displayed in Figure 1. Grouping based on $Overperform$ has a number of advantages. First, it allows us to reallocate capital across branches that under- versus over-perform. Second, it is orthogonal to observables (e.g., size, maturity, score); hence, reallocation does not have an incidence or disparate impact over demographics. Third, the reallocation gains are not due to optimal pricing of applicant and loan characteristics. Finally, averaging across geographies and observables allows us to average out shocks specific to the time of the experiment.

Calculations. The marginal return is a function of the slope of demand Θ' and cost of risk q'. We estimate these slopes using the instrumental variables approach described in Section 3.

The marginal return is also a function of the average return π and demand Θ, which are observed ex-post and can be estimated ex-ante. Capital allocation decisions are made ex-ante, whereas realizations include a noise or luck component. Our main calculations are based on *expected* values, using only information available to the bank.

For each bin b, we calculate the marginal return as the empirical counterpart of (Marginal):

$$\mu_b = \mathbb{E}[\pi_b] + \frac{\mathbb{E}[\Theta_b]}{\Theta'_b}(1 - q'_b). \tag{2}$$

We calculate the expectation $\mathbb{E}$ from the econometrician's perspective, taking into account all observables available to the bank, including *persistent unexpected overperformance*. This approach asks whether the bank allocates efficiently given the information it conditions on. Our main calculations are these expected values from the econometrician's perspective.

For completeness, we also report calculations based on realizations, replacing $\mathbb{E}[\pi_b]$ and $\mathbb{E}[\Theta_b]$ in (2) with the realized $\hat{\pi}_b$ and $\hat{\Theta}_b$. The marginal returns in expectation and realization are highly correlated, with a correlation coefficient of 0.94. Table 13 summarizes these calculations.

Finally, all bank-level statistics—the mean, the dispersion, and the reallocation gain—weight

[27] Balancing the two—the analogue of the optimal histogram rule of Scott (1979)—gives $B^* \sim N^{1/3}$. For N between 17,096 loans (cost of risk) and 50,042 applications (demand), $N^{1/3} \approx 26$–37; our choice of $B = 20$ sits just below this range and is therefore conservative.

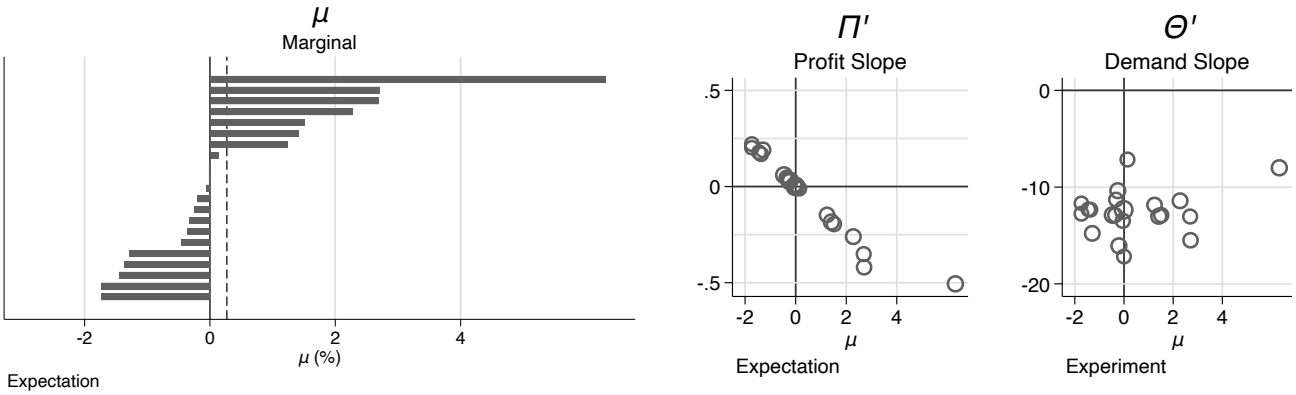
bins by their volume of lending in the year prior to the experiment. This matches the theory, where $\bar{\mu}$ is the volume-weighted mean satisfying the aggregate capital constraint.

Table 13: Results: Marginal Return Calculations—Expectation vs. Realization

	π	Θ	Θ'	q'
Expectation	$\mathbb{E}[\pi \mid X]$	$\mathbb{E}[\Theta \mid X]$	RCT	RCT
Econometrician Realization	$\hat{\pi}$	$\hat{\Theta}$	RCT	RCT

Note. $\mu = \pi + \frac{\Theta}{\Theta'}(1 - q')$. X is observables available to the econometrician, including *Overperform*. Levels are measured in the experimental baseline group. $\hat{\pi}$ and $\hat{\Theta}$ are realizations. Θ' and q' are estimated using the experiment.

Figure 2: Results: Marginal Returns



Note. The left panel displays the estimated μ for 21 bins, sorted in increasing order. The dashed vertical line indicates the weighted mean across all bins. The right panel displays Π' and Θ'—the effect of a $1pp$ APR change in R on Π and Θ, for each bin.

4.1 Marginal Returns

Figure 2 visualizes the marginal return calculations: the left panel displays the estimated marginal return for each bin, sorted in increasing order, and the right panel displays the underlying changes in profit and demand.

Demand is downward sloping for every bin, with an interquartile range of the slope from -13.0 to -11.7. The slope of profit, by contrast, takes on both positive and negative

values, with an interquartile range from -0.18 to 0.06. Because demand slopes are uniformly negative, the sign of the marginal return $\mu_b = \Pi'_b / \Theta'_b$ is determined by the sign of the profit slope.

Table 14 reports the estimates. We estimate the weighted mean of marginal returns to be 0.17%. The marginal return is highest at 6.21% for the online branch (i.e., customers acquired online, who are not assigned a branch manager).

We first test the null hypothesis that the mean marginal return is zero, $H_0 : \bar{\mu} = 0$. A zero mean is a necessary condition for unconstrained optimality, which requires $\mu_b = 0$ for every branch; when it fails, the bank could increase profits by moving toward the first-order condition. With a standard error of 0.30% and a p value of 0.58, the test for unconstrained optimality does not reject. Hence, we can not reject that the total scale of capital allocation is optimal. This finding also complements the institutional evidence that capital is not rationed: a binding aggregate constraint would imply a common *positive* marginal return, for which we find no evidence—though the confidence interval does not rule out a small one.

The test of whether the mean marginal return is equal to zero does not tell us whether the allocation is efficient. Allocative efficiency concerns the (lack of) dispersion in the marginal return. If one branch has a higher marginal return while another has a lower marginal return, then reallocating capital can increase overall profits.

Approximately 40% of branch groups (by lending volume) exhibit a positive marginal return, indicating that capital is under-allocated. As the demand curve is downward sloping, this implies the price should be lowered, which would increase the allocated capital. For the remaining branch groups, a negative marginal return indicates a need to decrease the allocated capital. This dispersion in marginal returns implies gains from reallocating capital to high-marginal-return branch groups, holding total allocation constant.

4.2 Reallocation

Next, we calculate the implied reallocation using (Reallocation).

Figure 3 visualizes the steps. The first figure displays the optimal price R^* vs the observed price $\hat{R}$, for each bin. The third figure displays the optimal allocation Θ^* vs the observed allocation $\hat{\Theta}$, for each bin. The second and fourth figures display the difference between the optimal and observed prices and allocations. By construction, the capital reallocated $\sum_b \Theta^*_b - \hat{\Theta}_b$ sums to zero.

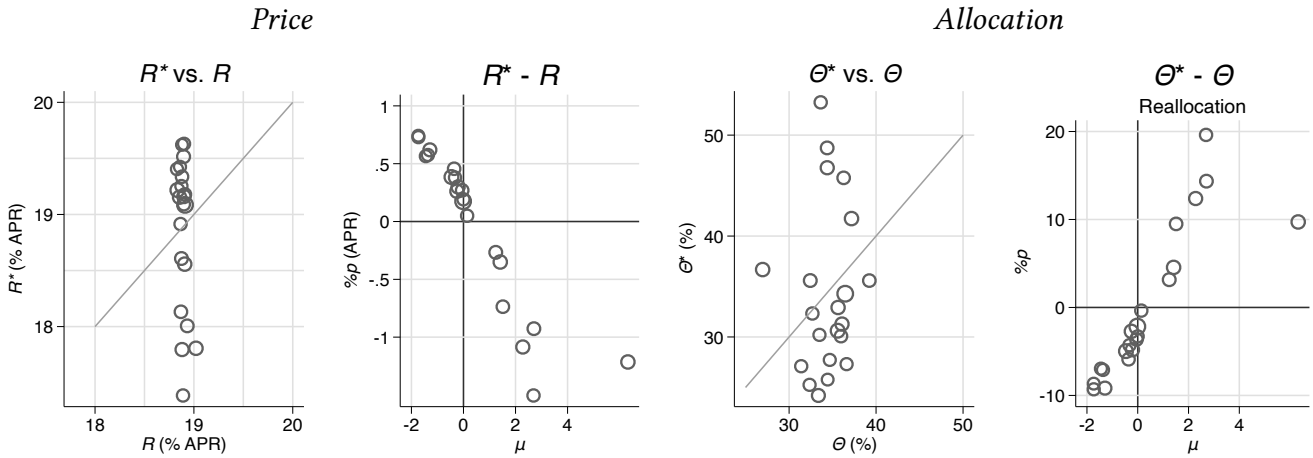
Table 14: Results: Marginal Return and Reallocation Gain

	μ **Marginal Return**		$\mathcal{G}$ **Reallocation Gain**		
	Estimate	p	Estimate	p	%
Panel A: Physical Branches					
Constrained			0.052 (0.015)	0.002	8.4
Unconstrained	0.17 (0.30)	0.58	0.055 (0.018)	0.007	8.8
Realization			0.059 (0.015)	0.001	9.4
Panel B: All Branches					
Constrained			0.065 (0.018)	0.002	10.7
Unconstrained	0.49 (0.43)	0.27	0.069 (0.022)	0.006	11.4
Panel C: Cross-sell and Origination Costs					
Constrained, Physical			0.052 (0.014)	0.001	8.4
Constrained, All			0.064 (0.017)	0.002	10.3
Panel D: Robustness—Winsorization					
Constrained, All			0.059 (0.015)	0.001	9.7

Note. Unless otherwise noted, all calculations using expected values from the econometrician's perspective, weighted (volume of lending in the year prior), and using reallocation gains formula (Gain). The first column reports the weighted mean of the bin-level marginal returns, $\bar{\mu}$, in % of face value. The third and fifth columns report the reallocation gain in % of face value per application and as a % of profits.

Directionally, a positive marginal return indicates that the price should decrease and the capital allocation should increase, as the second and fourth figures show. Quantitatively, the implied reallocation is larger when the marginal return is larger in abso-

Figure 3: Results: Optimal Price and Allocation



lute value, when demand is more responsive to price, and when the profit function is less curved; more sharply diminishing returns imply smaller reallocations for the same marginal return.

One striking pattern is the degree of uniformity in the observed prices, displayed in the first figure. Observed prices range from 18.83 to 19.02 (a 19 *bps* difference) whereas optimal prices range from 17.4 to 19.63 (a 223 *bps* difference). Thus, the optimal prices are an order of magnitude more dispersed than the observed prices, albeit mostly lying within the experimental price variation support (220 *bps*). These prices imply an allocation ranging from 24.2 to 53.3, displayed in the third and fourth figures.

4.3 Gain

Testing for *constrained* efficiency takes into account that the branches are connected through the internal capital market and share a common cost of funds c. A constrained efficient allocation features *marginal returns equal to a common value across branches* (Section 2.3): every branch then faces the same effective cost of capital, and no profitable reallocation across branches remains.

We calculate the implied gains using (Gain). The first row of Table 14 displays the results. We reject the null of allocative efficiency (p=0.002). Hence, there are benefits to reallocating capital across branches until marginal returns are equal. We find the potential gain from reallocation through equalizing marginal returns to be 0.052% of face value, corresponding to 8.4% of profits.

In the second row, we report the unconstrained gain, without the total capital constraint (i.e., move all branches to their first-order condition)—implying a gain of 8.8%.

Therefore, the reallocation gains arise primarily (8.4/8.8=95%) from the dispersion of marginal returns across branches, highlighting gains from reallocating capital while maintaining the total capital constraint. The remainder (5%) would come from expanding total lending by lowering the average price—though, as reported above, we cannot reject that the current scale is optimal.

4.4 Extensions

We provide several extensions of and robustness checks for our calculations. These results are reported in Panels B and C of Table 14.

Online branch. Next, we include the online branch, which accounts for about 5% of the applications during the experiment. At the online branch, applications arrive without a branch manager or loan officer: no in-person screening, no soft information, and less monitoring after origination. The channel also self-selects borrowers who avoid, or are not served by, the branch network, and its volume expanded during the pandemic, when in-branch contact was limited. The online branch has the lowest demand level and slope, and highest cost of risk level and slope.

Including the online branch, we calculate the reallocation gains to be higher, corresponding to 10.7% of total profits. Overall, we find that the mean marginal return is slightly larger than zero, but we cannot reject the null that the bank optimizes its first-order condition. However, we find that marginal returns display substantial dispersion across branch groups, and we reject equalized marginal returns (p=0.002), which points to benefits from reallocation.

Cross-sell and Issuance Costs. Our analysis so far ignores cross-sell and issuance costs. However, what might appear as an inefficient allocation of capital could represent a strategic decision to cross-sell. For example, a branch that issues loans with low returns may do so to retain customers who later purchase other products. This cross-subsidization could allow banks to optimize overall profitability, even if individual decisions and branches appear inefficient in isolation.

Relatedly, we have assumed the costs of issuing a loan consist only of the cost of funds c (which also includes regulatory costs) and the cost of risk q. However, we can also account for additional operating marginal costs associated with originating loans, such as application, origination, and underwriting costs at the branch (e.g., salaries, rent) and headquarter (e.g., technology) level. For example, some branches may have lower or higher processing costs because of economies of scale or labor and operational expenses.

We denote the profits from *cross-sell* of other products or services as o and the cost of origination by ζ. Then the average return π becomes:

$$R - c - q(R) + o - \zeta,$$

where o represents the present value of the profits from *cross-sell* of other products or services over the next year (i.e., the sum of income generated from checking and savings accounts; credit cards; as well as cross-sell of insurance and brokerage products). One year is also the horizon of the bank's own profitability attribution. The average cross-sell is 0.25% of FV, which corresponds to 14% of the net present value per dollar lent. Similarly, ζ represents the additional marginal costs associated with originating loans. We assume o does not change as price changes.

We measure these costs using data from the bank's internal operational expense accounting. This allows us to account for all costs in the same way the bank accounts for them. These marginal costs (about 0.02% of face value) are negligible in magnitude compared with the profit per dollar lent. Moreover, they vary minimally across branches (except for the online branch) and, therefore, have little effect on the constrained gain. The first row of Panel C reports these results, which are similar to our main estimates.

5 Discussion

We return to the puzzle posed in the introduction: an internal capital market, by virtue of its control rights, can reallocate capital across branches in ways that an external market cannot. Why, then, are marginal returns not equalized?

To begin, our partner bank exhibited *managerial inertia*—it was *not aware* of *persistent unexpected overperformance* at the branch level that we document in Section 1.2. The bank was attentive to branch performance: it monitored lending volume, prices and spreads, non-performing loans, and income generated for each branch-year. Its headquarters and regional managers were aware that in any given year branches may have higher or lower returns, but attributed such differences either to observables (e.g., credit score, size, maturity) or to transitory noise—an unlucky draw of defaults, a local shock, a one-off year that the next would reverse. Isolating the persistent component requires long-run data, pooling roughly 500 branches over several years: it is real but buried in annual noise, and becomes visible only in the pooled data. As a result, the bank based underwriting and pricing decisions solely on the credit score and did not adjust capital allocation in response to the *persistent unexpected* differences.

More broadly, our analysis indicates that the bank could achieve allocative efficiency by *reallocating* capital on the margin. Many levers could implement such reallocation (e.g., pricing, underwriting, incentives, discretion); a comprehensive analysis is beyond the scope of this paper. However, we show that prices can be used to equalize marginal returns—raising prices and redirecting capital away from low-marginal-return branches toward high-marginal-return branches.

Yet the bank priced *uniformly* across branches. Uniform pricing—while not a traditional friction—impedes reallocation: it prevents the bank from diverting capital away from branches with low marginal returns and toward branches with high marginal returns.[28]

Why was pricing uniform? Was the bank simply unable to price by branch (e.g., due to regulation)? In other contexts (e.g., retail) or countries (e.g., the U.S.), uniform pricing may be due to technological and institutional frictions. To assess the generalizability of our results, we briefly review these frictions.

Menu Costs. Uniform pricing is a rigidity, but across space. Traditionally, such rigidities are rationalized by menu costs. For example, in retail, changing shelf prices can be time and labor-intensive. In contrast, banks operate in a digital environment with real time pricing. Relatedly, banks may offer a *house rate* to all applicants or engage in *billboard pricing*—publicly advertise a price in brochures or on television—to attract applicants. Such practices are common among government banks, credit unions, and integrated lenders, yet at most banks (including our partner bank) applicants readily receive different prices based on their credit score.[29]

Regulation. Geographic pricing is common across retailers, banks, and insurance companies. Yet, in some countries and contexts, regulation may constrain geographic pricing.[30] In Türkiye, location is *not* a protected class (unlike sex, age, religion etc.), so

[28] Similar uniform pricing practices have been documented in the United States for retail chains (DellaVigna and Gentzkow, 2019), deposit pricing by banks (Granja and Paixao, 2026; Begenau and Stafford, 2023), and mortgages securitized by government-sponsored enterprises (Hurst et al., 2016). There, uniform pricing means the same posted price for the same good across locations; here, a common score-to-spread schedule. The reallocation gains we document are broadly comparable to those found in other industries. Most notably, DellaVigna and Gentzkow (2019) show that most U.S. food, drugstore, and mass-merchandise chains price uniformly across branches, resulting in a median store profit loss of 1.79% of *revenue*, which translates into 17.9% of profits at a 10% operating margin and 35% of profits at a 5% operating margin.

[29] As the lever is adjustable, dispersion in marginal returns cannot be rationalized as the shadow of adjustment costs on quasi-fixed inputs—a leading concern in the across-firm literature (Asker et al., 2014; David and Venkateswaran, 2019).

[30] In the U.S., for example, banks are known to offer lower deposit prices in remote and less competitive markets (e.g., North Dakota) compared to highly competitive urban markets (e.g., New York City); see Drechsler et al. (2017). Branch is an operational factor tied to service delivery (as opposed to the location

pricing on this basis complies with banking regulations.[31]

Securitization. Loans destined for securitization—particularly when government-sponsored—must conform to standardized criteria (e.g., credit score thresholds, loan-to-value ratios), and, crucially, standardized pricing structures.[32] Naturally, in our internal capital market setting, securitization is not relevant, as individual projects do not require external funding.

In our conversations with the bank, however, two concerns were repeatedly raised against branch-level pricing: strategic behavior and fairness.

First, bank managers worried that branch-level pricing could trigger *strategic* behavior: applicants might manipulate the branch of application to secure favorable terms. Low-price branches might attract disproportionate application volumes, undermining expertise and established relationships. In practice, however, the scope for such manipulation is limited: customers lack proximate alternatives; they are tied to, or can be directed to, a specific branch; and repeated or manipulated applications are flagged during underwriting.

The second concern raised is that branch-level pricing could create *perceptions of unfairness* among both customers and employees. Such perceptions may harm the brand and, in the long run, dampen employee morale, customer demand, and profits. From a customer point of view, however, the concern has limited purchase: the dimension along which prices would vary—persistent unexpected overperformance—is orthogonal to observables such as loan size, maturity, and credit score, and is not a proxy for any customer characteristic. The employee concern is harder to dismiss: branch staff and

of residence or collateral), so branch-level pricing constitutes segmentation, not discrimination. However, laws against redlining prohibit credit pricing based on the location of the asset (though not the applicant), including the Fair Housing Act, ECOA, and CRA. Nevertheless, U.S. firms frequently use zip codes to price other products such as auto insurance or supermarket vouchers. In France, by contrast, place of residence is explicitly protected under anti-discrimination law. Türkiye aligns more closely with the U.K. and Canada, neither of which recognizes location/place as a protected class in financial services.

[31] Reallocation using branch-level pricing is organizational: it conditions on an organizational characteristic, the branch, to offset an organizational friction. Our sorting variable—overperform—is orthogonal to size, maturity, score. Hence the reallocation holds applicant characteristics fixed, and therefore is not price discrimination—third-degree price discrimination conditions price on an applicant characteristic to exploit heterogeneous willingness to pay. Furthermore, reallocation and price discrimination differ in what the price varies with. Price discrimination varies prices with demand: across groups, price differences track differences in elasticities. Reallocative prices vary with returns: the price is the lever that moves capital, so branch prices differ whenever marginal returns do, even if demand and cost slopes are identical across branches. The price cut at a high-return branch does not exploit its demand; it directs capital toward it.

[32] In such cases, banks may originate loans to meet secondary market standards and optimize liquidity and risk transfer, effectively raising menu costs at the organizational level. This is documented for mortgages sponsored by government entities in the U.S. by Hurst et al. (2016).

managers are evaluated in part on volume, and those assigned less favorable terms may read the policy as an adverse judgment on their branch, or as handicapping them relative to their peers.

In sum, uniform pricing may reflect a response to these concerns rather than a failure to optimize: the relevant costs skew toward strategic and psychological frictions rather than institutional or technological ones. The bank can be thought of as maximizing $\Pi - \kappa$, where κ stands for the strategic and fairness considerations above and is incurred only under branch-level pricing. Uniform pricing is then optimal if and only if $\kappa \geq G^C$. Hence, the forgone reallocation gain is the shadow price of uniformity, and our estimate quantifies the threshold: the offsetting forces must amount to at least the gross gain for uniformity to be optimal.[33]

Alternative Levers. Our study focuses on a single, but arguably the most central, lever for reallocation: prices. Yet reallocation levers run the gamut (e.g., underwriting, volume targets, solicitation, and others). Each lever reallocates capital and earns the same π (by the product rule) by increasing lending on the extensive margin, yet has distinct inframarginal effects on those who take up. Hence, each lever carries its own reallocation cost, and therefore its own marginal return (and curvature).[34]

Furthermore, levers may also differ in salience (e.g., underwriting thresholds are invisible) and fairness: solicitation (inviting applicants at good branches) is unobjectionably fair; underwriting (same score, different access, and denial a greater harm than a higher price) arguably less so. Finally, levers may also affect the applicant pool (e.g., a lower branch price may attract disproportionately many applicants), a margin our experiment holds fixed and our gross gains do not price.

In Table 15 we discuss these levers in detail; yet analyzing all of them is beyond the scope of this paper, and we leave the empirical analysis for future work. Our findings are nonetheless interpretable as a lower bound, generally strict, on the multi-lever gain.

[33] Yet, other factors may cut in favor of branch-level pricing. To the extent that persistent overperformance reflects branch effort—screening and producing information (Stein, 2002; Liberti and Petersen, 2019)—the pricing regime determines whether that effort is rewarded. Uniform pricing severs this link: a branch that screens better faces the same rate and roughly the same book as one that does not.

[34] Formally, for any lever x that moves volume, the product rule gives

$$\Pi_x = \Theta_x \, \pi + \Theta \, \pi_x \qquad \Longrightarrow \qquad \mu_x = \frac{\Pi_x}{\Theta_x} = \pi + \frac{\Theta}{\Theta_x} \, \pi_x,$$

so the extensive-margin effect is the average return π, evaluated at the observed allocation and hence identical across levers, while the inframarginal term $(\Theta/\Theta_x)\pi_x$ is lever-specific, because each lever moves the average return along a different path. Marginal returns coincide across levers only when the bank sets every lever against a common internal shadow price.

6 Conclusion

We study whether a bank's internal capital market allocates capital efficiently across branches. Using a randomized pricing experiment, we identify the marginal return on capital for each branch group. We cannot reject that the total lending scale is optimal, but we reject that marginal returns are equalized across branch groups, implying gross reallocation gains of 8.4 to 10.7% of profits. We identify the bank's branch-uniform pricing policy as an important constraint on price-based reallocation: optimal prices would need to be far more dispersed across branches than the near-identical spreads the bank currently charges.

Our findings illustrate a tension at the heart of internal capital markets. These markets exist, in part, to exploit the advantages of centralized control over allocation. Yet the same headquarters may adopt uniform policies—whether for simplicity, fairness, or to deter strategic behavior—that prevent prices from serving as a reallocation tool. The gain we estimate is the gross profit gain from relaxing uniform pricing; whether the bank would fully realize that gain depends on implementation, strategic, and fairness costs that we do not observe. Nonetheless, this estimate provides a benchmark against which the costs of branch-level pricing can be assessed.

Our analysis highlights that average returns can be misleading about marginal returns and feasible reallocation gains. Reallocating capital through prices is costly because of inframarginal effects on existing loans, a cost that only experimental variation can measure directly. We do not claim that the gain magnitude generalizes beyond this bank; the broader contribution is about the mechanism. Uniform pricing is pervasive: it has been documented in U.S. retail chains (DellaVigna and Gentzkow, 2019), bank deposit markets (Granja and Paixao, 2026; Begenau and Stafford, 2023), and GSE-securitized mortgages (Hurst et al., 2016). Whether the allocative costs of uniformity documented here are representative of the broader banking sector—where deposit pricing, mortgage pricing, and cost-of-funds allocation also exhibit uniformity—is an important question for future work.[35]

More broadly, our approach applies whenever a decision-maker allocates a scarce resource across competing uses through an operational lever. The recipe requires only identified variation in the lever. The applications extend beyond prices or branches: al-

[35]Uniformity in the internal *cost of capital* is equally familiar. Graham and Harvey (2001) and Krüger et al. (2015) document that firms pervasively apply a single, firm-wide discount rate across divisions and projects — the *WACC fallacy*. Here uniformity is costly even absent differences in risk, so our findings generalize the fallacy: uniform internal prices of capital forgo profit whenever marginal returns differ, whatever their source.

locating capital across credit scores (Agarwal et al., 2018; Argyle et al., 2020); allocating a fixed advertising budget (Blake et al., 2015; Shapiro et al., 2021); tightening underwriting thresholds (Berg, 2015; DeFusco et al., 2020); and expanding approval margins (Karlan and Zinman, 2010). In each case there is a decision, an identified lever, and a forgone profit; the method turns the variation into a marginal return per use, and marginal returns into a gain.

References

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2018). Do banks pass through credit expansions to consumers who want to borrow? *Quarterly Journal of Economics 133*(1), 129–190.

Alchian, A. A. (1969). Corporate management and property rights. In H. G. Manne (Ed.), *Economic Policy and the Regulation of Corporate Securities*, pp. 337–360. Washington, DC: American Enterprise Institute.

Argyle, B., T. Nadauld, and C. Palmer (2020). Monthly payment targeting and the demand for maturity. *Review of Financial Studies 33*(11), 5416–5462.

Asker, J., A. Collard-Wexler, and J. De Loecker (2014). Dynamic inputs and resource (mis)allocation. *Journal of Political Economy 122*(5), 1013–1063.

Aydin, D. (2024). Forbearance vs. interest rates: Experimental tests of liquidity and strategic default triggers. Working paper.

Bau, N. and A. Matray (2023). Misallocation and capital market integration: Evidence from india. *Econometrica 91*(1), 67–106.

Begenau, J. and E. Stafford (2023). Uniform rate setting and the deposit channel. Working paper, Harvard Business School.

Berg, T. (2015). Playing the devil's advocate: The causal effect of risk management on loan quality. *Review of Financial Studies 28*(12), 3367–3406.

Bils, M., P. J. Klenow, and C. Ruane (2021). Misallocation or mismeasurement? *Journal of Monetary Economics 124*, 39–56.

Blake, T., C. Nosko, and S. Tadelis (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica 83*(1), 155–174.

Board of Governors of the Federal Reserve System (2026). Assets and liabilities of commercial banks in the united states (H.8 statistical release). Weekly release; accessed July 2026.

Campello, M. (2002). Internal capital markets in financial conglomerates: Evidence from small bank responses to monetary policy. *The Journal of Finance 57*(6), 2773–2805.

Carrillo, P., D. Donaldson, D. Pomeranz, and M. Singhal (2023). Misallocation in firm production: A nonparametric analysis using procurement lotteries. Working Paper 31311, National Bureau of Economic Research.

Cavallo, A., B. Neiman, and R. Rigobon (2014, 03). Currency unions, product introductions, and the real exchange rate*. *The Quarterly Journal of Economics 129*(2), 529–595.

Cortés, K. R. and P. E. Strahan (2017). Tracing out capital flows: How financially integrated banks respond to natural disasters. *Journal of Financial Economics 125*(1), 182–199.

Cremers, K. J. M., R. Huang, and Z. Sautner (2011, 10). Internal Capital Markets and Corporate Politics in a Banking Group. *The Review of Financial Studies 24*(2), 358–401.

Daruich, D. and J. Kozlowski (2023, July). Macroeconomic implications of uniform pricing. *American Economic Journal: Macroeconomics 15*(3), 64–108.

David, J. M., L. Schmid, and D. Zeke (2022). Risk-adjusted capital allocation and misallocation. *Journal of Financial Economics 145*(3), 684–705.

David, J. M. and V. Venkateswaran (2019). The sources of capital misallocation. *American Economic Review 109*(7), 2531–2567.

DeFusco, A. A., S. Johnson, and J. Mondragon (2020). Regulating household leverage. *Review of Economic Studies 87*(2), 914–958.

DeFusco, A. A., H. Tang, and C. Yannelis (2022). Measuring the welfare cost of asymmetric information in consumer credit markets. *Journal of Financial Economics 146*(3), 821–840.

DellaVigna, S. and M. Gentzkow (2019, 06). Uniform Pricing in U.S. Retail Chains*. *The Quarterly Journal of Economics 134*(4), 2011–2084.

Dlugosz, J., Y. K. Gam, R. Gopalan, and J. Skrastins (2024). Decision-making delegation in banks. *Management Science 70*(5), 3281–3301.

Dobbie, W., A. Liberman, D. Paravisini, and V. Pathania (2021). Measuring bias in consumer lending. *Review of Economic Studies 88*(6), 2799–2832.

Drechsler, I., A. Savov, and P. Schnabl (2017, 05). The Deposits Channel of Monetary Policy*. *The Quarterly Journal of Economics 132*(4), 1819–1876.

Duchin, R. and D. Sosyura (2013). Divisional managers and internal capital markets. *The Journal of Finance 68*(2), 387–429.

Einav, L. and A. Finkelstein (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives 25*(1), 115–138.

Erickson, T. and T. M. Whited (2000). Measurement error and the relationship between investment and q. *Journal of Political Economy 108*(5), 1027–1057.

Froot, K. A. and J. C. Stein (1998). Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics 47*(1), 55–82.

Gertner, R. H., D. S. Scharfstein, and J. C. Stein (1994, 11). Internal versus External Capital Markets*. *The Quarterly Journal of Economics 109*(4), 1211–1230.

Gilje, E. P., E. Loutskina, and P. E. Strahan (2016). Exporting liquidity: Branch banking and financial integration. *The Journal of Finance 71*(3), 1159–1184.

Giroud, X., G. Matvos, A. Seru, and R. Silva (2022). Managing resource (mis) allocation.

Giroud, X. and H. M. Mueller (2015). Capital and labor reallocation within firms. *The Journal of Finance 70*(4), 1767–1804.

Gopinath, G., Ş. Kalemli-Özcan, L. Karabarbounis, and C. Villegas-Sanchez (2017). Capital allocation and productivity in south europe. *The Quarterly Journal of Economics 132*(4), 1915–1967.

Graham, J. R. and C. R. Harvey (2001). The theory and practice of corporate finance: Evidence from the field. *Journal of Financial Economics 60*(2–3), 187–243.

Granja, J. and N. Paixao (2026). Bank consolidation and uniform pricing. *Journal of Financial Economics 176*.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica 50*(1), 213–224.

Hazell, J., C. Patterson, H. Sarsons, and B. Taska (2025). National wage setting. Working Paper 30623, National Bureau of Economic Research.

Hitsch, G. J., A. Hortaçsu, and X. Lin (2021). Prices and promotions in U.S. retail markets. *Quantitative Marketing and Economics 19*(3), 289–368.

Houston, J., C. James, and D. Marcus (1997). Capital market frictions and the role of internal capital markets in banking. *Journal of Financial Economics 46*(2), 135–164.

Hsieh, C.-T. and P. J. Klenow (2009, 11). Misallocation and Manufacturing TFP in China and India*. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Hughes, D. and J. Majerovitz (2024). Measuring misallocation with experiments. Working paper.

Hurst, E., B. J. Keys, A. Seru, and J. Vavra (2016, October). Regional redistribution through the us mortgage market. *American Economic Review 106*(10), 2982–3028.

Jayaratne, J. and P. E. Strahan (1996). The finance-growth nexus: Evidence from bank branch deregulation. *The Quarterly Journal of Economics 111*(3), 639–670.

Karlan, D. and J. Zinman (2010). Expanding credit access: Using randomized supply decisions to estimate the impacts. *Review of Financial Studies 23*(1), 433–464.

Krüger, P., A. Landier, and D. Thesmar (2015). The wacc fallacy: The real effects of using a unique discount rate. *The Journal of Finance 70*(3), 1253–1285.

Lamont, O. (1997). Cash flow and investment: Evidence from internal capital markets. *The Journal of Finance 52*(1), 83–109.

Landier, A., D. Sraer, and D. Thesmar (2017). Banking integration and house price comovement. *Journal of Financial Economics 125*(1), 1–25.

Laufer, S. and A. Paciorek (2022). The effects of mortgage credit availability: Evidence from minimum credit score lending rules. *American Economic Journal: Macroeconomics*. Verify year, volume, pages.

Lewis, R. A. and J. M. Rao (2015). The unfavorable economics of measuring the returns to advertising. *Quarterly Journal of Economics 130*(4), 1941–1973.

Liberti, J. M. and M. A. Petersen (2019). Information: Hard and soft. *Review of Corporate Finance Studies 8*(1), 1–41.

Maksimovic, V. and G. Phillips (2002). Do conglomerate firms allocate resources inefficiently across industries? theory and evidence. *The Journal of Finance 57*(2), 721–767.

Matvos, G. and A. Seru (2014, 02). Resource allocation within firms and financial market dislocation: Evidence from diversified conglomerates. *The Review of Financial Studies 27*(4), 1143–1189.

Nakamura, E. (2008, May). Pass-through in retail and wholesale. *American Economic Review 98*(2), 430–37.

Park, K. and G. Pennacchi (2009, 05). Harming Depositors and Helping Borrowers: The Disparate Impact of Bank Consolidation. *The Review of Financial Studies 22*(1), 1–40.

Peek, J. and E. S. Rosengren (1997). The international transmission of financial shocks: The case of japan. *The American Economic Review*, 495–505.

Scharfstein, D. S. and J. C. Stein (2000). The dark side of internal capital markets: Divisional rent-seeking and inefficient investment. *The Journal of Finance 55*(6), 2537–2564.

Scott, D. W. (1979). On optimal and data-based histograms. *Biometrika 66*(3), 605–610.

Shapiro, B. T., G. J. Hitsch, and A. E. Tuchman (2021). Tv advertising effectiveness and profitability: Generalizable results from 288 brands. *Econometrica 89*(4), 1855–1879.

Shin, H.-H. and R. M. Stulz (1998, 05). Are Internal capital Markets Efficient?*. *The Quarterly Journal of Economics 113*(2), 531–552.

Sraer, D. and D. Thesmar (2023, April). How to use natural experiments to estimate misallocation. *American Economic Review 113*(4), 906–38.

Stein, J. C. (1997). Internal capital markets and the competition for corporate resources. *The Journal of Finance 52*(1), 111–133.

Stein, J. C. (2002). Information production and capital allocation: Decentralized versus hierarchical firms. *The Journal of Finance 57*(5), 1891–1921.

Stein, J. C. (2003). Agency, information, and corporate investment. Handbook of the Economics of Finance.

Weitzman, M. L. (1974, 10). Prices vs. quantities. *The Review of Economic Studies 41*(4), 477–491.

Whited, T. M. (2001). Is it inefficient investment that causes the diversification discount? *The Journal of Finance 56*(5), 1667–1691.

Whited, T. M. and J. Zhao (2021). The misallocation of finance. *The Journal of Finance 76*(5), 2359–2407.

Williamson, O. E. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.