Uniform Pricing and Capital Allocation *

Deniz Aydın MIT

Ernest Liu Princeton

Janis Skrastins Utah

January 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 reallocated. We cannot reject the null that total lending scale is optimal, but we reject equalized marginal returns across branch groups (p=0.002): marginal returns are substantially dispersed, implying gross potential reallocation gains of 10.7% of profits. We show that the bank's branch-uniform pricing policy—charging the same spread across all branches conditional on credit score—is an important friction that limits price-based reallocation.

*We are grateful to David Sraer for extended discussions; to Cevdet Akcay, Philip Dybvig, Ali Hortacsu, David Thesmar, Pete Klenow, Gyongyi Loranth, Jeremy Majerovitz, Martin Oehmke, Antoinette Schoar, Adi Sunderam; 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 Spring 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) for valuable suggestions. Washington University IRB (HRPO #202209011, *Capital Reallocation*).

Introduction

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, 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.[1] 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 reallocated via price changes. We cannot reject the null that total lending scale is optimal (p=0.27), but we reject equalized marginal returns across branch groups (p=0.002): marginal returns are substantially dispersed, implying gross potential reallocation gains of 10.7% of profits. 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), Hughes and Majerovitz (2024)). 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 both more directly testable and more actionable: the question is not whether a social planner could improve on decentralized markets, but whether the firm's own allocation mechanism can be improved.

A notable feature of our analysis is that *average* returns are misleading about the scope for reallocation. Using historical data, we document that the average return at top-decile branches is nearly four times that at bottom-decile branches. This raises concern about misallocation, but does not by itself measure feasible gains from reallo-

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

cation. 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. We define the marginal return as:

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

where $\Pi'(R)$ and $\Theta'(R)$ are the responses of profit and demand (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 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. This variation allows us to observe what happens when capital is diverted to or away from a branch, and to estimate the curvature of the profit function without imposing a common functional form across branches.

In Section 4, we present the findings.[2] We estimate the weighted mean of marginal returns across branch groups to be 0.49% (s.e. 0.43%), and we cannot reject unconstrained optimality (p=0.27). However, we reject equalized marginal returns across branch groups (p=0.002). Approximately 60% of branch groups (by lending volume) exhibit a positive marginal return, suggesting that capital should increase, while the remainder exhibit negative marginal returns. The implied gross reallocation gain is 10.7% of profits. These gains arise predominantly (94%) from the dispersion of marginal

[2]We group the bank's approximately 500 branches into 21 branch groups based on persistent unexpected overperformance to ensure sufficient statistical power within each group. All references to "branch groups" refer to these groupings; see Appendix C for details.

returns across branches, highlighting the value of *reallocating* existing capital rather than expanding total lending. The remaining 6% could be realized by increasing total lending. Our results are robust to accounting for cross-sell, origination costs, alternative binning, and winsorization.

It may be surprising that capital allocation is inefficient *within* an organization that, by virtue of its control rights, can reallocate in ways that external markets cannot. 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 can be interpreted as the total cost of accommodating these concerns. Notably, after reviewing our results, the bank began revising its practices in favor of branch-level pricing.

Literature. Our paper contributes to the literature on the efficiency of internal capital markets in finance.[3] 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. (2010); Gilje et al. (2016); Cortés and Strahan (2017); Drechsler et al. (2017)), but evidence on whether they allocate capital *efficiently* has been limited.[4] We contribute by conducting a field experiment to facilitate capital reallocation within a bank, test whether marginal returns are equalized, and quantify 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), 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

[3]See Lamont (1997); Shin and Stulz (1998); Gertner et al. (1994); Stein (1997); Scharfstein and Stein (2000); Whited (2001); Duchin and Sosyura (2013). 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) decompose misallocation across and within firms and find that three-quarters is within.

[4]Our paper also relates to financial integration more broadly, e.g., Jayaratne and Strahan (1996); Landier et al. (2017).

procurement lotteries in Ecuador, and Hughes and Majerovitz (2024) use randomized microcredit 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., 2017; DellaVigna and Gentzkow, 2019; Daruich and Kozlowski, 2023) for retail and (Park and Pennacchi, 2008; Granja and Paixao, 2021; Dlugosz et al., 2023) for banks—as well as studies featuring randomized pricing experiments.[5] 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 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. Borrowers commonly use these loans for debt consolidation, home improvements, health expenses, travel, education, or large purchases.

Underwriting. At the time of application, applicants provide basic demographic information (e.g., education, employment, home ownership, income). They then declare

[5]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, resulting in a median profit loss of 1.79% of *revenue*, which translates to 35% of profits at a 5% operating margin. Granja and Paixao (2021) 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.

Table 1: Environment: Underwriting

Policy rule Application rejected if larger than	Tier:	Prime	Mid	Sub
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.

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 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.

Internal Capital Market. The bank comprises $\approx$ 500 branches. These branches face distinctive pools of quality and competition, informational access, and some decision-making autonomy. The branches access a common pool of capital at a common cost of funds; hence, the opportunity cost of the capital deployed in one branch is a loan in

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

Score	Cost of Funds	Spread	Rate
	% APR	% APR	% 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.

another branch. There are no lending or deposit constraints at the branch level, and the movement of capital across branches is unrestricted.[6][7] [8]

[6]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.

[7]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.

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

1.2 Returns to Capital—Measurement and Dispersion

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

1.2.1 Measurement

Data. We analyze historic 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.[9]

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 dispersion. Loans originating from the worst-performing branches have substantially lower returns than those originating from the best-performing branches.

Focusing on 2015-2020, the branch-level return range (10th to 90th percentile) is $1.45pp$ 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 realized loan-level realized returns

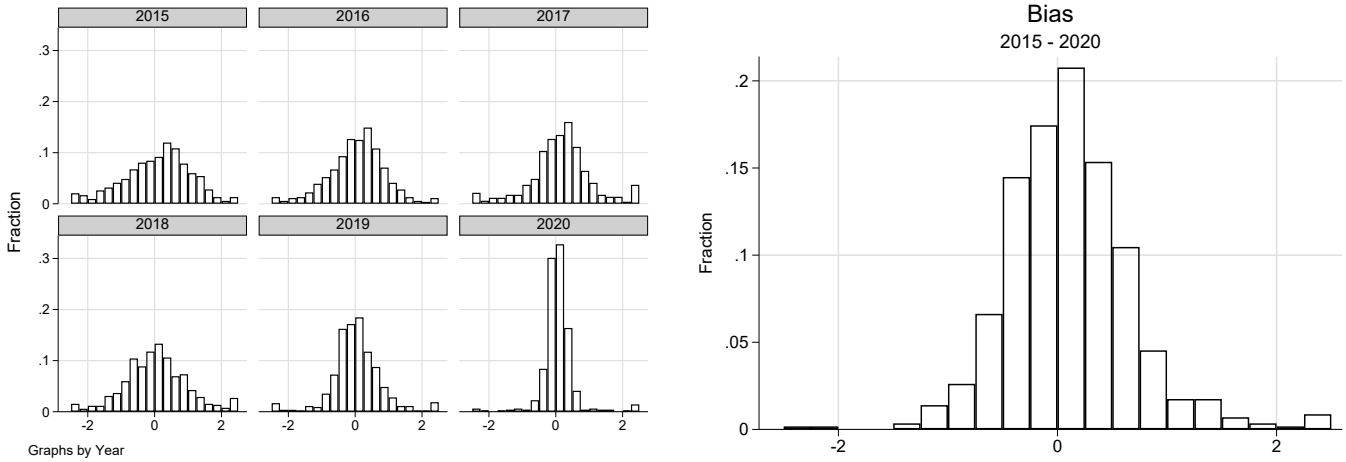
[9] Further details are provided in Appendix A.1.

π_j^R of observables X_j.[10]

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

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.

Figure 1: Historic Data: Average Returns—Branch Overperformance



Note. The left figure displays the histogram of branch-year effects f_{bt} from Equation (1). 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*—as the time average of these branch-year effects. The right panel of Figure 1 displays the histogram of this branch-level bias. $Overperform$ gives us a natural running variable for our reallocation exercise in Section 4. Indeed,

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.

[10] X consists of credit score, loan size, and maturity. 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.

2 Conceptual Framework

We now describe our approach to identifying marginal returns using supply variation, and linking the marginal return estimates to tests of allocative efficiency and 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).[11] Here, we focus on reallocation across branches using pricing variation.

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 and $q_b(R_b)$ the cost of risk (missed payments per dollar lent). We assume q is increasing in R—reflecting adverse selection and moral hazard.

As in Section 1, the average return (i.e., *average profit per capital allocated*) is:[12]

$$\pi_b(R_b) = R_b - c - q_b(R_b) \qquad \text{(Average)}$$

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:

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

[11] See Table 12.

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

2.1 Unconstrained Optimality

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:[13]

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

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. Because the demand curve is downward-sloping, allocating more capital requires lowering prices. Lowering prices increases demand, where each new loan brings additional profit—an additional average return— (and vice versa for raising prices). The second term represents the intensive margin—the *inframarginal* effect. Lowering prices affects the returns on loans already taken up, through changes in price and the cost of risk. Optimality is achieved by balancing these two effects.

2.2 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:

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

[13] While the approach we develop use and the experiment creates pricing variation focus on pricing, they are inherently tied to the capital allocation problem. Deploying more capital requires lowering the price because the demand curve is strictly downward-sloping. Indeed, holding features of the technology (e.g., underwriting, advertising) constant, the marginal μ we identify is equivalent to the marginal of the capital allocation problem against an inverse demand. See early work by Weitzman (1974) and Table 12 for a discussion of this equivalence.

μ rationalizes the observed price $\hat{R}$ and may be interpreted as a *misperception* about cost of funds c.[14]

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

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

The change in profit at $\hat{R}$, equals the product of μ and the slope of the demand. Hence, μ can be estimated, using price variation around $\hat{R}$, as:

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

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'.

Unconstrained optimality (with no total capital constraint) requires the chosen price satisfies the first-order condition, corresponding to the null hypothesis $H_0 : \mu = 0$. When $\mu \neq 0$, an unconstrained bank can increase profits. 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).

2.3 Constrained Optimality

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

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

[14] 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.

The prices are then chosen as:

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

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 constraint.[15] When 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 (i.e., holding total allocation constant) across uses, we derive an analytical formula. The bank's reallocation gains due to the *dispersion* in marginal returns in the presence of constraint is:

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

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 $(\hat{\mu}_b - \bar{\mu})$, 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

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

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

prices such that μ is common across uses at the level $\bar{\mu}$.

The reallocation gains 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 is 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.[16]

2.5 Taking Stock

We define an allocative inefficiency as a situation where the bank, given its constraints and information, could use price adjustments to reallocate total capital and increase overall profits. This definition of efficiency requires that the marginal returns are equalized across uses, though this common marginal return need not be zero. If one use has a higher marginal return (indicating that allocating more capital would increase profits) and vice versa, then reallocation increases profits. To the extent that μ_b is not constant across uses, the bank can improve profitability by adjusting prices to reduce or eliminate marginal return dispersion.

Table 4: Conceptual Framework: Average, Marginal, Reallocation, Gain

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

Table 4 summarizes the concepts used throughout the paper—average return, marginal return, reallocation, and gain.

[16] Under linear demand and cost curves (i.e., $\Theta'' = q'' = 0$), the gain simplifies to $\frac{1}{4} \, \mu^2 \, \frac{1}{(1-q')}$.

Table 5: 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%	-10	-10	1	5	-10	-40	10%
1%	-10	-20	-0.5	-2.5	5	20	5%

Note. $\hat{\Pi} = 0.5\%$, $\hat{R} = 20\%$ APR, $\hat{\Theta} = 25\%$.

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 the μ and the difference of the observed price $\hat{R}$ from the optimal level R^*:

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

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

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

Hence, the unconstrained reallocation gain is given by:

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

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

$$\mathcal{G}^U_{\text{Total}} = \sum_b \mathcal{G}^U_b(\hat{\mu}_b)$$

The analytical formula for the reallocation gains of a *constrained bank* follows analogously. It only remains to be shown that constraint is satisfied.

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

$$\begin{aligned}\Theta^* - \Theta(\hat{R}) &\approx \Theta'(\hat{R})\big(R^* - \hat{R}\big)\\ &\approx\ -\hat{\mu}\,\frac{\Theta'(\hat{R})^2}{\Pi''(\hat{R})}\end{aligned}$$

In constrained reallocation, we replace $\hat{\mu}_b$ with $\hat{\mu}_b - \bar{\mu}$:

$$\Theta^*_b - \Theta_b(\hat{R}_b) \approx\ -\big(\hat{\mu}_b - \bar{\mu}\big)\,\frac{\Theta'_b(\hat{R}_b)^2}{\Pi''_b(\hat{R}_b)} \qquad \text{(Reallocation)}$$

The total change in capital allocation is thus $\sum_b(\bar{\mu} - \hat{\mu}_b)\ \frac{\Theta'_b(\hat{R}_b)^2}{\Pi''_b(\hat{R}_b)}$, 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 heterogeneous marginal returns at more granular levels.

Design. The experiment involves 50,042 applicants approved for underwriting.[17] 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%) receives 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. 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.[18]

Summary Statistics. Table 6 displays summary statistics. Panel A displays applicant characteristics. Panel B displays loan characteristics. The average price, cost of funds, and spread is 18.8%, 15.7%, and 3.2%. The average maturity is 29 months and the average requested face value is 29,611 TRY, about six months of income. The maximum maturity is three years, with two-thirds of all loans having a maturity of three years.

First, the applicants decide whether to take-up (demand). A total of 17,096 applications (34.2%) were taken up and issued. Later, those who take-up, decide whether to miss payments (cost of risk). Among takers, 3.7% fall into 90+ (delinquent) status, resulting in a cost of risk of 1.72% of the face value. This yields an average realized

[17] The experiment ran from November 2020 to January 2021. It was conducted by the bank's treasury, with the goal of estimating the average slope of demand and profit curves. 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.

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

Table 6: 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
Cost of Origination	ζ	% of FV	50,042	0.07	0.22	0	0	0.24
Panel C: Outcomes								
Demand	Θ	%	50,042	34.2	47.4	0	0	100
90+		%	17,096	3.7	18.9	0	0	0
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

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.

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, ma-

Table 7: Experiment: Covariate Balance

	Metro	**Score**	**Income**	**Loans**	**Payments**	**Debt**	**Term**	**Cost of Funds**
TRY	Count	TRY	TRY	Months	% APR			
$\mathbb{Z}^L$	-0.003	21.8	232	-0.01	-8.7	633	-0.3	0.003
Low	(0.007)	(3.6)	(162)	(0.03)	(28)	(624)	(0.14)	(0.01)
$\mathbb{Z}^H$	-0.012	-15.2	43	-0.02	-13.9	441	-0.2	0.04
High	(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

turity) as fixed.[19] 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.

Table 8: Experiment: Estimation Framework

	Z	**X**	**Y**
	Instrument		
First-stage	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	R
Reduced Form	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	Θ or q
Treatment Effect	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	R	Θ or q

[19] Table 13 shows that a small fraction 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 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 vary across experimental groups.

Estimation Framework. Table 8 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 effect 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.[20]

Table 9: 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
$\mathbb{Z}^L$	-1.32	-1.32	13.2	-0.68
Low	(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)
$\mathbb{Z}^H$	0.92	0.88	-16.6	0.41
High	(0.01)	(0.01)	(0.7)	(0.35)
N	50,042	50,042	50,042	17,096
F	7,286	12,603		

Average Effect Estimates. Table 9 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).

[20]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.

Focusing on the first-stage estimates in Table 9 the baseline group receives an 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 9 show a downward sloping demand curve. The average demand for the baseline group is 34.5%, which is $13.2pp$ (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%). This yields a convex demand, with an slope of -12.7.

The second column of reduced form estimates in Table 9 show 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.

4 Results

We now operationalize the framework developed in Section 2.

Table 10: 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

Table 10 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 -1.43%—= $\frac{34.2\%}{-12.7}(1-0.47)$—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 is close to optimal.

Yet, these estimates mask substantial 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 $B \approx 500$ branches. To estimate marginal returns and reallocate across, we group loans into bins. Binning ensures a sufficient number of observations per bin and reduces statistical uncertainty in estimated demand and cost functions.[21]

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 (for π and Θ), based on the *same information* as the bank. For completeness, we also report calculations based on realizations.[22]

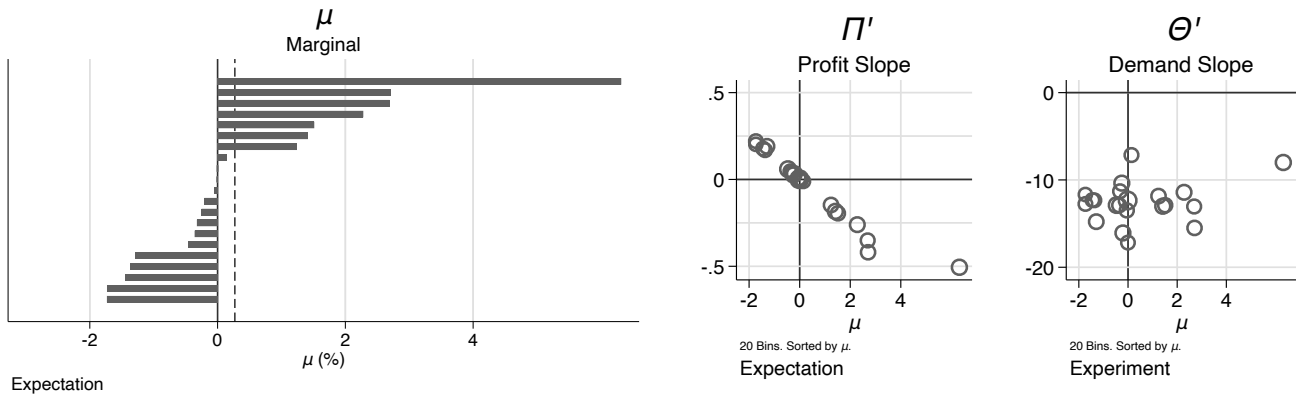
4.1 Marginal Returns

Figure 2 visualize the marginal return calculations. The marginal return is the change in profit per marginal capital allocated. The right panel displays the change in profit and change in demand separately. The left panel displays the estimated marginal returns

[21] 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 gains are not due to *optimal pricing* of observables. Finally, averaging across geographies and observables allows us to average out shocks specific to the time of the experiment. See Appendix C for details.

[22] We calculate the expected return two ways: expectation (econometrician) and expectation (bank). Both expectations use the observables available to the bank. The econometrician, in addition to observables, also takes into account *persistent unexpected overperformance*, whereas the bank does not. These approaches answer different questions: is the bank allocating efficiently given the information it conditions on versus is the bank allocating efficiently given the information it could condition on but does not? Our main calculations are expected values from the econometrician's perspective. The three marginal returns we estimate are highly correlated, with a pairwise correlation of 0.99, 0.94, and 0.89. Figure 5 displays the expectation and realizations for demand Θ and cost of risk q.

Figure 2: Results: Marginal Returns



Note. The left panel displays the estimated μ for 20 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.

for each bin, sorted in increasing order. The dashed vertical line indicates the weighted mean across all bins.

Both the demand and cost slope exhibit substantial variation. Demand is downward sloping for every bin, with an interquartile range of -13.0 to -11.7. The slope of the profit takes on both positive and negative values, with an interquartile range of -0.18 to 0.06. Hence, the sign of the marginal return is determined by the sign of the slope of the profit function.

Table 11 reports the estimates. We estimate the weighted mean of marginal returns to be 0.49%. The marginal return is highest at 6.21% for the online branch (i.e., customers acquired online and are not assigned a branch manager).

We first test the null hypothesis that the mean marginal return is equal to zero, or $H_0 : \mu_b = 0$—*degeneracy at the value of zero for all branches.* This condition is generically necessary for unconstrained optimality, as the bank could increase profits by moving to the first order condition. With a standard error of 0.43% and a p value of 0.27, the test for unconstrained optimality does not reject.[23] Hence, we cannot reject that the total scale of capital allocation is optimal.

[23]How do marginal returns relate to average returns? First, marginals take on both positive and negative values. Second, Recall that the μ equals the average return π plus a term representing how costly it is to reallocate (that depends on demand Θ, demand slope Θ' and cost of risk slope q').

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 60% 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: Results: Optimal Price and Allocation

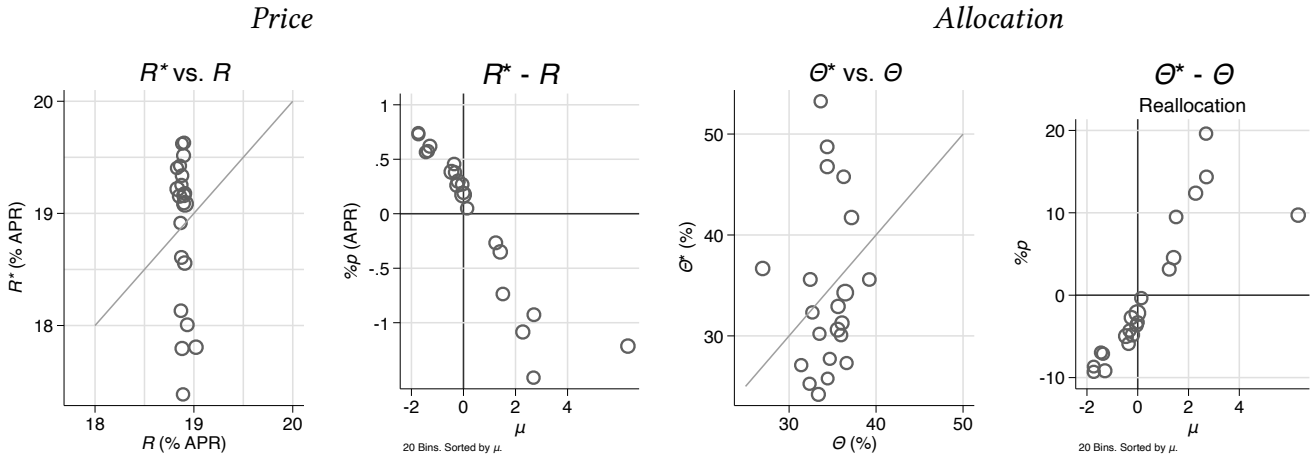


Figure 3 visualizes the steps. The first figure displays the optimal price R^* vs the observed price $\hat{R}$, for each bin. The third figure displays the optimal allocation Θ^* vs the observed allocation $\hat{\Theta}$, for each bin. The second and fourth figures display the difference between the optimal and observed prices and allocations.

Directionally, a positive marginal return indicates that price should decrease and capital allocation should increase. This can be seen in the second and fourth figures. Quantitatively, the magnitude of reallocation also depends on the slope of demand

and inversely related to the curvature of the profit function. Hence, reallocation is higher when the absolute value of the marginal return is higher but also when demand slope is higher and curvature of the profit function is lower. Naturally, more sharply diminishing returns to scale and smaller reallocations for the same marginal return. We estimate curvature using finite differences after estimating the slope separately for price cuts and hikes.

One striking pattern is the degree of uniformity in the observed prices, displayed in the first figure. Observed prices range from 18.89 to 19.02 (a 13 bps difference) whereas optimal prices ranges from 17.39 to 19.63 (a 224 bps difference). Thus, the optimal price is about *20 times* more dispersed compared with the observed prices. These prices imply an allocation ranging from 24.2 to 53.3, displayed in the third and fourth figures. This dispersion in optimal allocations is three times greater than that of observed allocations.

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. Hence, a constrained efficient allocation would feature *marginal return equal to a common value across branches*, though this common value is not necessarily zero. A common marginal return ensures that all branches are operating under the same pricing distortions, eliminating inefficiencies that arise from variation in pricing distortions.

We calculate the implied gains using (Gain). The first row of 11 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 gain from potential reallocation through equalizing marginal returns to be 0.065% of face value, corresponding to 10.7% of average returns.

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 11.4%. Therefore, the reallocation gains arise primarily (10.7/11.4=94%) from the dispersion of marginal returns across branches, highlighting gains from reallocating capital while maintaining the total capital constraint. However, the positive mean marginal return indicates further gains (6%) by decreasing the average price across all branches.

Table 11: Results: Marginal Return and Reallocation Gain

	μ **Marginal Return**		$\mathcal{G}$ **Reallocation Gain**		
	Estimate	p	Estimate	p	%
Panel A: Main Estimates					
Baseline	0.49 (0.43)	0.27	-0.065 (0.018)	0.002	-10.7
Unconstrained			-0.069 (0.022)	0.006	-11.4
Realization	0.36 (0.35)	0.32	-0.062 (0.014)	0.0003	-10.3
Panel B: Extensions					
Cross-sell and Cost of Origination	0.65 (0.42)	0.14	-0.064 (0.017)	0.002	-10.5
Excluding Online	0.97 (0.52)	0.08	-0.052 (0.015)	0.002	-8.4
Panel C: Robustness					
Expectation—Bank	0.63 (0.50)	0.22	-0.076 (0.027)	0.01	-11.9
Bin by Size	0.56 (0.47)	0.24	-0.082 (0.019)	0.0004	13.7

Note. Unless otherwise noted, all calculations based on B=21 bins by $Overperform$, using expected values from the econometrician's perspective, weighted (volume of lending in the year prior), and using reallocation gains formula (Gain).

4.4 Extensions and Robustness

We provide several extensions and robustness checks and for our calculations. These results are reported in Panels B and C of Table 11.

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). Similarly, ζ represents the additional marginal costs associated with originating loans.

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 effects on constrained loss. The third row of Panel C reports these results. We find results similar to our main estimates.

Online branch. Next, we exclude the online branch, which account for about 5% of the applications during the experiment. The online branch has the lowest demand level and slope, and highest cost of risk level and slope. Excluding the online branch, we calculate the reallocation gains to be slightly lower—0.052% of face value, corresponding to 8.4% of total profits.

Robustness. We also provide a number of robustness checks. First, instead of grouping branches into 20 bins in our baseline estimation, we use 10 bins.[24] Second, we group

[24]The dispersion in marginal returns could also reflect statistical uncertainty of the estimates. How-

branches by size. Third, as the quadratic approximation to the loss is sensitive to outliers, as a robustness check, we winsorize marginal returns at 5% level. We find broadly consistent results.

Overall, we find that the mean marginal return is slightly larger than zero, but we cannot reject the null that banks optimize their 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.

5 Discussion

It may be surprising that a single organization fails to allocate capital efficiently *within*. After all, an internal capital market, by virtue of its control rights, can actively reallocate capital across branches in ways that an *external* market cannot. Why, then, is the allocation of capital inefficient?

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 understood that branch-year returns can differ, but attributed such differences either to observables (e.g., credit score) or to transitory noise. As a result, underwriting and pricing decisions were based solely on the credit score, and capital allocation was not adjusted 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. The potential reallocation levers (e.g., pricing, underwriting, incentives, discretion) run the gamut, and a comprehensive analysis is beyond the scope of this article. However, our analysis indicates 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, in our setting, the bank was engaging in pricing that is *uniform* across branches. When pricing is uniform across branches—while not a traditional friction—it acts as an impediment to reallocation, preventing diversion of capital away (to) from projects with lower (higher) marginal return.[25]

ever, in our setting with 50,000, we have on average $\approx$ 1,000 observations per bin, and hence the standard errors of our estimates are an order of magnitude smaller than the observed cross-sectional dispersion in marginal returns. Moreover, when we coarsen the groups, the pattern of dispersion remains robust.

[25] Similar uniform pricing practices have been documented in the United States for retail chains

Why was pricing uniform, in particular, was the bank unable (e.g., due to regulations)? In other contexts (e.g., retail) or countries (e.g., U.S.), uniform pricing may be due to technological and institutional frictions. To discuss the generalizability of our results, we briefly discuss 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.

Regulation. Geographic price discrimination is common across retailers, banks, and insurance companies. Yet, in some countries and contexts, regulation may constrain geographic pricing.[26] In Türkiye, location is *not* a protected class (unlike sex, age, religion etc.), so pricing on this basis is not non-compliant with banking regulations.

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.[27] Naturally, in our internal capital market setting, securitization is not relevant, as individual projects do not require external funding.

(DellaVigna and Gentzkow, 2019), deposit pricing by banks (Granja and Paixao, 2021; Begenau and Stafford, 2023), and mortgages securitized by government-sponsored enterprises (Hurst et al., 2016). The reallocation gains we document are broadly comparable to those found in this literature 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.

[26] 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 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 recognize location/place as a protected class in financial services.

[27] 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).

However, in our conversations, two concerns were frequently raised as to why uniform pricing—strategic behavior and fairness.

First, it was indicated that branch-level pricing could trigger strategic behavior, as applicants manipulate application branches 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; are tied to, or can be directed to, a specific branch; with repeated or manipulated applications 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, dampen employee morale and customer demand (and profits) in the long run.

Notably, after reviewing our results, the bank began revising its practices in favor of branch-level pricing.

In sum, uniform pricing may be due to costs that skew more toward strategic and psychological frictions rather than institutional or technological ones. Our potential reallocation gains could be interpreted as the total costs of these offsetting forces.

6 Conclusion

We study whether a bank's internal capital market allocates capital efficiently across branches. Using a randomized pricing experiment, we identify marginal returns—the change in profit per dollar of capital reallocated via price changes—for each branch group. We cannot reject that total lending scale is optimal, but we reject that marginal returns are equalized across branch groups, implying gross reallocation gains of 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 10.7% 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.

More broadly, 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 10.7% 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, 2021; 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.

References

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

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.

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, Donaldson, Pomeranz, and Singhal (2023). Misallocation in firm production: A nonparametric analysis using procurement lotteries.

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 (2010, 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.

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 (2023). Decision-making delegation in banks. *Management Science 0*(0), null.

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.

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.

Granja, J. and N. Paixao (2021). Market concentration and uniform pricing: Evidence from bank mergers. Technical report, Bank of Canada Staff Working Paper.

Hitsch, G., A. Hortacsu, and X. Lin (2017). Prices and promotions in us retail markets.

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.

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.

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 co-movement. *Journal of Financial Economics 125*(1), 1–25.

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 (2008, 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.

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.

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*.

Appendices

A Returns to Capital

A.1 Calculations

The data contains information on the calendar date the loan has first reach 30+, 60+, 90+ and non-performing status; as well as recovery made in the 12 month consequent to defaults for non-performing loans.

First, we calculate a base NPV for loans that *did not* default. Second, for loans that *did* default, we calculate a base present value assuming uniform recovery over 12 months subsequent to the NPL date.

$$NPV = -FV + \frac{\text{Pay}}{c}\left(1 - \frac{1}{(1+c)^T}\right) \tag{2}$$

$$NPV = -FV + \frac{\text{Pay}}{c}\left(1 - \frac{1}{(1+c)^{T^{NPL}-1}}\right) + \frac{1}{(1+c)^{T^{NPL}}}\left(\frac{\frac{\text{Rec}}{12}}{c}\left(1 - \frac{1}{(1+c)^{12}}\right)\right) \tag{3}$$

Finally, we make adjustments for late payments. Suppose a single payment Pay due at a period t was instead made at period $t' > t$. In that case, the present value of the loan would need to be adjusted:

$$\begin{aligned} PV' - PV &= \text{Pay}\left(\frac{1}{(1+c)^{t'}} - \frac{1}{(1+c)^t}\right) \\ &= \frac{\text{Pay}}{(1+c)^t}\left(\frac{1}{(1+c)^{t'-t}} - 1\right) \end{aligned}$$

We use this formula to make adjustments, considering all possible scenarios.

B Conceptual Framework: Additional Results

B.1 Prices vs. Allocations

C Binning

Table 14 displays summary statistics for the bins and shows minimal variation in observables across bins. For example, the interquartile range ($p75$–$p25$) of credit score, loan size, and maturity across bins is a mere 1%, 3%, and 2% of the respective averages. Figure 4 displays the average observables by bin, sorted by estimated μ in increasing order, further demonstrating that these observables are uncorrelated with the estimated marginal returns. For instance, the correlation between marginal return and credit score is -0.16 ($p = 0.26$). We also provide *robustness checks* in which we residualize, and our results remain quantitatively consistent. Figure 4 corroborates that any variation across bins is indeed orthogonal to marginal returns. As a robustness check, we also conduct reallocations based on alternative criteria, such as branch size.

Figure 4: Applicant and Loan Characteristics, by Bin

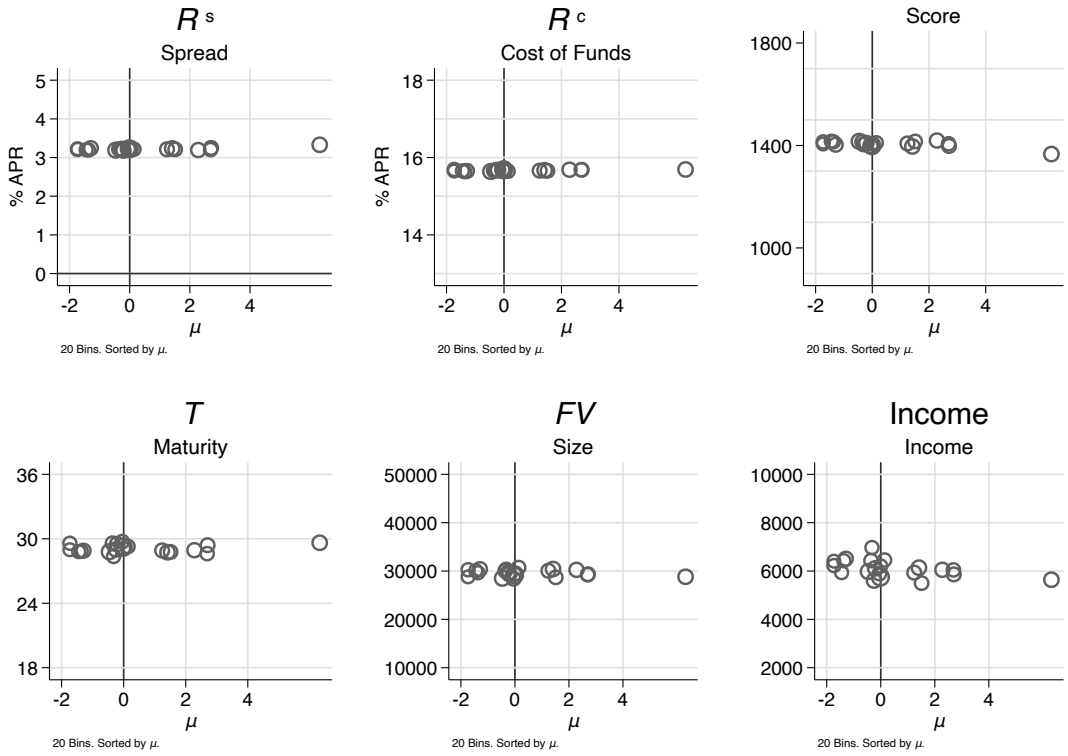
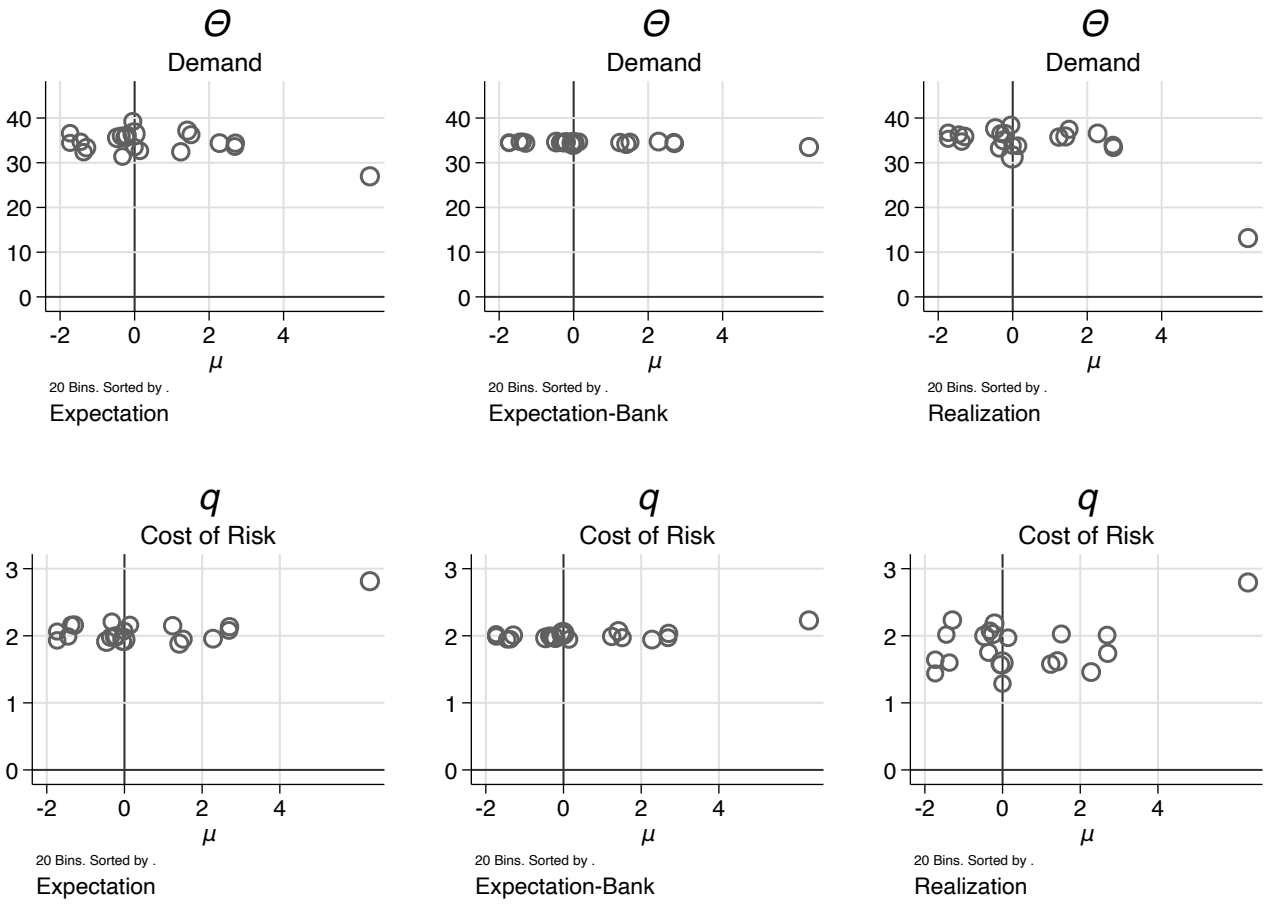


Figure 5: Marginal Return: Expectation vs. Realization



Note. Figure shows Top row plots Estimated μ for 20 bins, sorted in an increasing manner. Vertical line is the weighted mean.

Table 12: Conceptual Framework: Prices vs. Allocations

Prices

$$R^* = \arg\max_R \Theta(R)\Big(R - c - q(R)\Big) \qquad \Pi_R(R^*) = 0$$

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

$$\Pi_R(R) = \Theta_R(R)\,\Big(R - c - q(R)\Big) + \Theta(R)\,\Big(1 - q_R(R)\Big)$$

$$\mu = \pi + \tfrac{\Theta}{\Theta_R}\,(1 - q_R)$$

Allocations

$$\Theta^* = \arg\max_\Theta \Theta\,\Big(R(\Theta) - c - q(R(\Theta))\Big) \qquad \Pi_\Theta(\Theta^*) = 0$$

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

$$\Pi_\Theta(\Theta) = \Big(R(\Theta) - c - q(R(\Theta))\Big) + \Theta\Big(R_\Theta(\Theta) - q_R(R(\Theta))\, R_\Theta(\Theta)\Big)$$

$$\mu = \pi + \Theta\big(R_\Theta - q_R R_\Theta\big)$$

Underwriting

$$u^* = \arg\max_u \Theta(u)\Big(R - c - q(u)\Big) \qquad \Pi_u(u^*) = 0$$

$$\hat{u} = \arg\max_u \Theta(u)\Big(R - c - q(u) - \mu\Big) \qquad \Pi_u(\hat{u}) = \mu\, \Theta_u(\hat{u})$$

$$\Pi_u(u) = \Theta_u(u)\Big(R - c - q(u)\Big) - \Theta(u) q_u(u)$$

$$\mu = \pi - \tfrac{\Theta}{\Theta_u} q_u$$

Solicitation

$$d^* = \arg\max_d \Theta(d)\Big(R - c - q(d)\Big) \qquad \Pi_d(d^*) = 0$$

$$\hat{d} = \arg\max_d \Theta(d)\Big(R - c - q(d) - \mu\Big) \qquad \Pi_d(\hat{d}) = \mu\, \Theta_d(\hat{d})$$

$$\Pi_d(d) = \Theta_d(d)\Big(R - c - q(d)\Big) - \Theta(d) q_d(d)$$

$$\mu = \pi - \tfrac{\Theta}{\Theta_d} q_d$$

Note. In the top panel, the bank chooses price R to maximize $\Pi(R)$ given demand $\Theta(R)$ and cost of risk $q(R)$. In the bottom panel, chooses an allocation K to maximize $\Pi(K)$ given inverse demand $R(K)$ and cost of risk $q(R(K))$.

Table 13: Effects on Other Margins

	T **Term**				*FV* **Size**			
	Changer Fraction %		Change %		Changer Fraction %		Change %	
	$\Delta > 10\%$	$\Delta > 1\%$	$\Delta > 10\%$	$\Delta > 1\%$				
$\mathbb{Z}^L$	0.3	-0.1	-2.1	-1.1	-0.2	-0.6	-0.2	-2.4
Low	(0.5)	(0.6)	(6.3)	(4.8)	(0.5)	(0.5)	(1.9)	(2.1)
$\mathbb{Z}^H$	0.3	-0.1	8.2	7.3	-0.7	-0.9	-4.1	-3.9
High	(0.8)	(1.0)	(9.8)	(7.5)	(0.7)	(0.8)	(3.2)	(3.4)
Baseline	5.9	8.1	111.7	107.6	4.8	5.8	58.7	65.0
	(0.2)	(0.2)	(2.4)	(1.8)	(0.2)	(0.2)	(0.7)	(0.8)
N	17,096	17,096	1,022	1,386	17,096	17,096	813	976
p	0.86	0.97	0.64	0.59	0.62	0.28	0.42	0.29

Table 14: Results: Summary Statistics

Variable		**Mean**	**SD**	**p25**	**p75**	**IQR**	**%**
Score (Credit Bureau)		1,406	11.7	1,399	1,414	15.1	1
Income		6,079	351	5,862	6,408	546	9
Size	FV	29,586	698	29,230	30,252	1,022	3
Maturity	T	29.1	0.37	28.8	29.4	0.60	2
Cost of Funds	c	15.7	0.02	15.7	15.7	0.03	0.2
Spread	$R - c$	3.22	0.03	3.2	3.2	0.02	0.7
Price	R	18.9	0.04	18.9	18.9	0.04	0.2
Demand	Θ	34.5	2.6	33.4	36.3	2.9	8
Slope	Θ'	-12.5	2.3	-13.0	-11.7	1.4	11
Curvature	Θ''	-362	233	-512	-134	378	104
Cost of Risk	q	2.07	0.21	1.94	2.15	0.21	10
Slope	q'	0.55	0.70	0.12	0.88	0.76	139
Curvature	q''	44.8	92.1	0.0	42.9	42.9	96
Profit	Π	0.61	0.09	0.55	0.67	0.12	20
Slope	Π'	-0.05	0.21	-0.18	0.06	0.24	505
Curvature	Π''	-30.1	7.8	-39.3	-20.2	19.0	63
Average	π^E	1.69	0.17	1.61	1.83	0.22	13
Marginal	μ^E	0.49	1.9	-0.46	1.41	1.88	386
Price (Optimal)	R^*	18.9	.66	18.6	19.3	0.78	4.1
Allocation (Optimal)	Θ^*	34.6	8.1	27.7	36.7	9.0	26
Reallocation	$\Theta^*_b - \hat{\Theta}_b$	0.11	8.3	-5.9	4.6	10.4	

Note. Table displays summary statistics for the B=21 bins. Θ, q, Π, π, and μ are expected values from the econometrician's perspective. Slope and curvature of Θ and q are estimated using the experiment. See Section 4 for details.