Uniform Pricing and Capital Allocation*

Deniz Aydın WUSTL
Ernest Liu Princeton
Janis Skrastins WUSTL

March 2025

Abstract

A fundamental question in finance is whether capital is allocated within a financial institution to equalize marginal returns across divisions. Yet, numerous sand-in-the-gears can lead to inefficiencies, while other practices may hinder the necessary reallocation. In this paper, we develop new techniques to test for *allocative efficiency* and quantify the resulting profit losses. These methods use pricing variation to identify divisional marginal returns, requiring minimal assumptions about functional forms, sources of inefficiencies, and the extent of optimizing behavior. We apply these techniques to study a large bank's internal capital markets through a field experiment that facilitates reallocation by supplying capital at prices assigned randomly. We statistically reject efficiency, and estimate a profit loss of 17%. Uniform pricing across divisions, while not a traditional friction and not necessarily the root cause, emerges as the key impediment to efficiency: it prevents banks from conditioning on geography and hence pricing from being used as a tool for capital reallocation to address frictions that lead to divisional underperformance. This lack of contingency in pricing has redistributive effects: subsidizing lower-income borrowers and divisions that persistently underperform expectations. We discuss implications for inequality and policy.

*We are grateful to David Sraer for extended discussions; to Cevdet Akcay, Suleyman Alan, Philip Dybvig, Ali Hortacsu, David Thesmar, Pete Klenow, Jeremy Majerovitz and Niccolo Ubertalli; and seminar participants at the Olin Business School, Koc University, American Economic Association meetings, U.S. Treasury (Office of Financial Research) for valuable suggestions.

A fundamental question in finance is whether capital is allocated within a financial institution to equalize marginal returns across projects and divisions. Yet, numerous *sand-in-the-gears* (e.g., incentives, information, authority, competition) can create inefficiencies, while practices and regulations (e.g., uniform pricing, redlining) may hinder necessary reallocation. Measuring the impact of these distortions and quantifying the resulting profit losses remains a significant challenge. Empirical tests of this hypothesis are scarce, with experimental evidence virtually nonexistent.

In this paper, we provide novel tests of efficiency in the pricing and allocation of capital within a financial institution. The two cornerstones of our investigation are: first, a novel methodology that uses pricing variation to test for efficiency and quantify the resulting profit losses; and second, a high-stakes, large-scale randomized natural field experiment at a large European bank that facilitates capital reallocation through pricing shocks. These pricing shocks, by inducing more (or less) capital allocation, enable us to identify the marginal returns and apply our methodology.

We apply this framework to study the efficiency of bank internal capital markets that allocate deposits across branches at a shared cost of funds. Bank internal markets, intermediating tens of trillions in the U.S., are arguably the most significant yet remain understudied compared to industrials and conglomerates.[1] Our study breaks new ground by conducting the first natural field experiment with a randomized control to test the efficiency of capital allocation within a financial institution.[2]

We proceed in five steps. First, we begin the empirical analysis by introducing our data and economic environment in Section 1. Using historical data and natural variation, we document large and persistent dispersions in the return to capital across organizational divisions. For instance, over a 6-year period, the realized profit per dollar lent at the 10th and 90th percentiles differ by 3% per dollar lent, nearly twice as much as the average return on capital. This evidence *invites the interpretation* that capital may be inefficiently allocated, with the bank leaving money on the table by allocating too much capital to low-return divisions, and vice versa. However, this analysis reflects *average* returns and does not directly reveal the *marginal* (next best) returns when capital is reallocated to different divisions.

Second, in Section 2, we develop a novel methodology that uses pricing variation

[1] See St. Louis Fed. The capital allocation processes within banks are understudied largely due to the inaccessibility of internal data and the lack of variation necessary for exogenous reallocation.

[2] The setting has unique advantages. The experiment analyzes application-level data and measure returns comprehensively (e.g., including timing and exposure at default, cross-sell, and origination costs) using administrative data and without measurement error. As we examine a single internal capital market and a single homogeneous product, inefficiencies are particularly puzzling. Crucially, our study's distinguishing feature is the experimental reallocation of capital, identifying the marginal return by underwriting loans that would not otherwise be issued.

to test for efficiency and quantify the resulting profit losses. An econometrician observes divisions operating within a financial institution, each offering loans under a given initial pricing structure. We say an *allocative inefficiency* exists if the bank (given constraints and information) could use prices to induce a reallocation of capital in a way that increases profits. Testing for efficiency would be straightforward if marginal returns—how a division's profits change as it lends more (or less) as prices changes—were directly observable. However, marginal returns are unobservable and differ in unknown ways. Additionally, failing to account for heterogeneity, nonlinearities, or returns to scale could be mistaken for inefficiencies. Our experimental approach circumvents these challenges by avoiding restrictive assumptions about functional forms, the extent of optimizing behavior, or the underlying sources of inefficiencies.

Our approach for testing and quantifying shares many commonalities with the wedge accounting framework. In the misallocation literature, the wedge represents a gap between the price of an input and its marginal value product. In our context, the wedges represents a distortion in the pricing first-order condition. The wedge is the treatment effect of a price change on profit per application, capturing both marginal and inframarginal effects, and can be identified using exogenous variation in price:

$$\mu = \frac{\sigma - 1}{\sigma} R - \frac{\sigma - \delta}{\sigma} q - c \qquad \text{(Wedge)}$$

where the expected profit per application—a function of interest rate R, cost of risk q, and demand and cost sensitivities σ and δ—is weighted off against to the cost of funds c. A positive wedge indicates that the bank should increase the price (implying a reduction in the allocated capital) and vice versa. Each division's wedge can be identified based on treatment effects on demand and costs. As a result of these wedges, the bank prices (and allocates capital) differently than what would be optimal.

We define two notions of efficiency. Unconstrained efficiency tests the optimality of divisional decisions in isolation, requiring that wedges are equal to zero for each division—i.e., each division is on its first-order condition. Constrained efficiency, on the other hand, evaluates the potential benefits of reallocation under an aggregate constraint, recognizing that divisions operate within a shared internal capital market and a common cost of funding. It requires that *wedges equal a common value across* divisions. As a result of these wedges, the bank prices (and hence allocates) capital differently than what would be optimal.

Our uses this logic, using the distribution of treatment effects from exogenous price variation to infer the degree of profit loss. Distinctively, we provide an analytical formula to express profit losses as a function of the cross-sectional dispersion of wedges

across divisions:

$$\mathcal{L}_{\text{Bank}} = \sum_b \mathcal{L}(\mu_b) = -\sum_b \frac{\mu_b^2}{2}\xi_b \qquad \text{(Loss, Total)}$$

While our novel approach can be applied to many questions and settings involving pricing, we focus on the allocative efficiency of lending across the divisions of a financial institution.

In Section 3, to obtain exogenous variation, we analyze a randomized field experiment involving $\approx$ 50,000 loan applications across $\approx$ 500 divisions. At status quo, the bank employed *risk-based* pricing using the public credit score that was *uniform* (i.e., constant markup on top of cost of funds and risk) across divisions. The randomized pricing experiment gives us a counterfactual world in which prices were determined randomly, at the loan level, either lower or higher than the baseline. Raising prices diverts capital away from a division, whereas lowering prices channels capital to a division.

We trace the response of capital deployed and profits generated in response to these random shocks by merging administrative data on take-up and realized performance. In response to a $\%p$ APR decrease in prices (from a baseline 18.9% APR), demand increases by $10\%p$ (off a base of 34.5%), increasing allocated capital by 29%. Conversely, in response to a $\%p$ APR increase in prices, demand decreases by $18\%p$, decreasing allocated capital by 52%. The changes in prices also generate an inframarginal effect—influencing the returns on loans already taken up—through changes in prices and the cost of risk (e.g., adverse selection). These average estimates mask substantial division-level heterogeneity. By comparing these responses across divisions, our approach enables us to test for inefficiencies, gauge the benefits of reallocation, and quantify the resulting profit losses.[3]

Fourth, in Section 4, we present the findings. First, we test for unconstrained efficiency—optimality of divisional decisions in isolation with no quantity constraint on lending— corresponding to the null hypothesis that all wedges are equal to zero. The average treatment effect of a supply-driven price decrease on profit per application is different from zero. Specifically, we estimate that the sales-weighted mean of wedges across divisions is 0.87 (standard error: 0.49), with $p = 0.09$. Hence the test for unconstrained efficiency does not reject. Approximately half of the divisions exhibit a negative wedge, suggesting the division is overpricing—the price should be lowered, which would increase the allocated capital. For the remaining divisions, a positive wedge

[3]The techniques we develop use pricing variation but directly study capital allocation, as the two are equivalent: allocating more capital—without altering origination quality—requires lowering the price due to the strictly downward-sloping demand curve. See early work by Weitzman (1974).

is indicating the division is underpricing, and hence a need to decrease the allocated allocated capital.

The test of unconstrained efficiency does not quantify how much money is left on the table. We then turn to *constrained efficiency*—gains from a reallocation across divisions holding lending constant, i.e., *factor-augmenting* reallocation—corresponding to a test of the null hypothesis that wedges take a common value across all divisions. Our findings reveal that wedges are dispersed, with a variance of 1.4 and an interquartile range of 1.9. The null hypothesis of constrained efficiency is decisively rejected ($p = 0.001$). We find that large profit gains from potential reallocation. Using a second-order approximation of profit losses, enabled by the experiment's inclusion of both price cuts and hikes, we estimate profit losses to be 17.5%. These losses arise from two primary sources. First, and importantly, the dispersion of wedges across divisions highlights potential gains from reallocating capital while maintaining the aggregate amount deployed. Second, the positive average wedge indicates further gains could be realized by increasing the average price across all divisions. Our results remain quantitatively consistent when we take into account cross-sell of other products as well as the additional marginal costs associated with originating loans.

Finally, in Section 5, we discuss why efficiency is not achieved and explore broader implications. It may be surprising that capital allocation is inefficient *within an institution* (as opposed to *across institutions*). After all, a single internal capital market does not face the commonly invoked financial frictions and by virtue of its control rights, can actively reallocate capital across divisions in ways that an *external* market cannot. Moreover, even if a friction (e.g., incentives, information, authority, competition) causes a low return to capital at some divisions, the price mechanism could still be used to reallocate capital away from low-return divisions to high-return divisions. After all, efficiency only requires prices to be adjusted to eliminate wedges. Why aren't divisions with a negative wedge not allocated more capital?

We argue that banks fail to efficiently allocate capital across divisions because they engage in risk-based pricing that is *uniform* across divisions. In other words, banks set a single price across all divisions based on observable risk, without accounting for *unobservable* characteristics and organizational frictions that cause some divisions to persistently underperform expectations. To highlight the quantitative importance of the uniformity: we find that the variation in demand sensitivity and cost of risk across divisions is approximately *100 times* greater than that of the price. Hence, uniform pricing, while not the root cause of inefficiency and not a traditional friction, acts as an impediment that prevents prices from being used as a tool for capital allocation, hindering the necessary reallocation of resources within the internal capital market to offset divi-

sional underperformance. To achieve optimal capital allocation across divisions, banks should adopt flexible pricing strategies that account for organizational frictions that the credit score does not capture, as well as local demand and cost.

Exactly identical prices hard to explain except by some kind of friction. Why do banks fail to account for variation in division-level factors in their pricing decisions? Various explanations, such as traditional rigidities due to technological and menu costs, billboard pricing, and securitization is not relevant in our context. The main driver of this practice is *regulatory and legal constraints*; most notably, *place-based discrimination* regulations which prohibit practices that deny individuals in certain areas equitable access to credit through discriminatory pricing or underwriting. Banks (and credit scores) must comply with various competition, anti-discrimination, consumer protection, and housing laws that collectively regulate practices that could be exclusionary. Divisions may face varying costs that credit scores fail to capture—stemming from unobservable factors or organizational frictions. However, regulations restrict financial institutions from conditioning pricing on geographic information. In Appendix C, we comprehensively discuss these laws. If allowed, region-specific risk factors could improve capital allocation. By being unable to adjust for regional differences, institutions are effectively forced to treat all divisions similarly. Our experimental design randomly relaxes this uniform pricing practice, and quantifies the economic costs of these regulatory and legal constraints.

We conclude by discussing several implications. Most notably, uniform pricing has redistributive effects, creating winners and losers: first, by underpricing and hence subsidizing lower-income borrowers at the expense of higher-income borrowers; second, by subsidizing divisions that persistently underperform expectations—a redistribution pattern that mirrors the dynamics of socialistic internal capital markets.

Literature. Our analysis and methods are inspired by the literature in macroeconomics that studies the misallocation of capital *across* firms (e.g., Hsieh and Klenow (2009)). Gopinath et al. (2017) documents dispersion of the return to capital across firms in Southern Europe. Bau and Matray (2023) use the staggered liberalization of access to foreign capital across industries to show that financial market changes can reduce misallocation. This literature often relies on naturally occurring data and highly parametric assumptions to infer marginal product of capital.[4] Two notable exceptions are Carrillo et al. (2023), who use construction procurement lotteries in Ecuador, and Hughes and Majerovitz (2024), who use randomized microcredit studies from Sri Lanka.

[4] Also see Banerjee and Moll (2010); David and Venkateswaran (2019); Sraer and Thesmar (2023). However, such procedures may yield biased estimates in the presence of endogeneity, misspecification, and heterogeneity.

In finance, prior studies have used naturally occurring data to document the existence and test the efficiency of internal capital markets within conglomerates and non-financial firms.[5] Similarly, the banking literature documents that bank internal capital markets channel deposits across divisions (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)).[6] Our study breaks new ground in this literature, by conducting the first high-stakes natural field experiment with a randomized control to facilitate capital reallocation and test whether marginal returns across divisions are equalized within an organization. The effective subsidy of underperforming divisions we document is akin to findings of Scharfstein and Stein (2000).

Our study also contributes to the literature on pricing, particularly on lack of contingency in general, and uniform pricing practices across divisions of large firms (Nakamura, 2008; Cavallo et al., 2014; Hitsch et al., 2017; DellaVigna and Gentzkow, 2019; Daruich and Kozlowski, 2023) and banks (Park and Pennacchi, 2008; Granja and Paixao, 2021; Dlugosz et al., 2023), as well as studies featuring randomized pricing experiments.[7] Cavallo et al. (2014) shows Zara prices are identical across countries within the Euro area, despite large differences in income and competition. DellaVigna and Gentzkow (2019) show that most U.S. food, drugstore, and mass-merchandise chains price uniformly across divisions, resulting in a median store 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 predominately use uniform *deposit* rate setting policies. Hurst et al. (2016) provide evidence of price uniformity for mortgages securitized by government-sponsored enterprises (GSEs), but not for prices in the private market. To the best of our knowledge, this is the first paper on uniform *loan* pricing in banks. We contribute to this literature by formalizing a general method through which exogenous pricing shocks can be employed to test for optimality and quantify profit losses. Our experiment creates a deviation from the risk-based uniform pricing strategy and finds that banks could increase profits by 17% if they *could* adopt geographically differentiated pricing strategies.

[5] 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) studies how conglomerates allocate resources across divisions over the business cycle. Matvos and Seru (2014) estimates a structural model of internal capital markets to study resource allocation within in response to external capital market dislocation. Giroud and Mueller (2015) studies capital reallocation within firms—using new airline routes as exogenous shocks to investment opportunities, documenting reallocation within, but only if the firm is financially constrained. Giroud et al. (2022) decompose misallocation across and within firms and find that three-quarter of the reallocation gains is due to misallocation within.

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

[7] e.g., DeFusco et al. (2022), Aydin (2024)

1 Environment and Data

1.1 Institutional Details

Here, we describe the key institutional features, including the collaborating institution, the product studied, underwriting, pricing, and incentives.

Collaborator. For the study we collaborate with a large French bank that is incorporated, operated, and regulated in the E.U. Our analysis uses data exclusively from the bank's subsidiary in Turkey. This subsidiary ranks among the top 10 in terms of asset size and number of divisions in the local market, with a borrower base representative of the local banked population. It provides a wide range of financial products, including unsecured loans, mortgages, checking and savings accounts, overdrafts, retirement accounts, brokerage, and insurance services.

Unsecured loans. We focus on a single product—unsecured loans with fixed interest rates, terms of up to 72 months, and fixed nominal installment payments in local currency. These loans account for two-thirds of the total non-mortgage *FV* outstanding to the household sector. Borrowers typically seek this credit for purposes such as debt consolidation, home improvements, health expenses, travel, education, or large purchases (e.g., mobile phones, tablets, and computers).

Table 1: Underwriting (Stylized)

Policy rule Application rejected if larger than	Tier	Prime	Mid	Sub
Debt-to-income		10	8	6
Payment-to-income		50%	40%	30%
Number of applications		10	10	5

Underwriting. At the time of application, applicants provide basic demographic information, including education level, employment type and title, home ownership status, and monthly disposable income. They then indicate the amount they wish to borrow, select a loan term, and declare a reason for borrowing.

Table 1 outlines a highly stylized underwriting policy. Applicants are categorized into tiers based on their credit scores. These tiers determine debt-to-income and payment-to-income limits. Underwriting is segmented such that salary customers and retirees, whose income is directly deposited into bank accounts and therefore verifiable, are subject to separate underwriting policies.

Underwriting allows for limited managerial discretion. Division managers may be asked to review and approve marginal applicants, but they cannot modify the interest rate, maximum loan term, or approved loan amount. Division managers, primarily incentivized based on lending quantity, play a key role in soliciting loan applications to increase sales volume.

Pricing. The interest rate is determined by adding a spread to the cost of funds. Pricing is *risk-based*: the spread depends solely on the credit score, which serves as a proxy for the cost of risk (i.e., the present value of missed payments). All experiment participants, excluding a small fraction ($\approx 10\%$) who receive a loyalty rate (e.g., bank employees, salary customers, retirees), are assigned one of five spreads. Table 2 shows that these five credit score groups alone explain more than 98% of the variation in the spreads. As credit scores are the sole determinant of pricing, interest rates (i.e., markups beyond the cost of funds and cost of risk) are effectively *uniform* across divisions.

Table 2: Risk-based Pricing

5 Credit Score Groups	✓		✓	✓
Branch		✓	✓	
Controls				✓
R^2	0.983	0.018	0.983	0.986

Note. Experiment participants in the baseline ($\mathbb{Z} = 0$) and excluding *loyalty rate* ($\approx$ 10%).

Internal Capital Market. We study an internal capital market, akin to those examined in the conglomerate literature. The bank comprises approximately 500 organizational divisions that share a common cost of funds. The divisions face distinctive pools of credit quality, competitive landscapes, informational access, and some decision-making autonomy. The bank has a long list of mutually exclusive applications where the opportunity cost of the deployed capital is a loan in another division. Bank finances projects centrally, from its own resources, or from a pool of capital (equity, debt, or deposits) raised externally. Hence, individual projects do not need individual external funding. Capital is (softly) budgeted at the start of the fiscal year through targets set at the division level, as well as across segments and products. Divisions are incentivized to achieve these targets. However, beyond these targets, there are no financial or funding constraints at the division level and the movement of capital across divisions is unrestricted. Additionally, there are no explicit costs associated with allocating excess capital to any particular division.

1.2 Returns to Capital—Measurement and Heterogeneity

Here, we introduce our data and the measures used to assess returns to capital in our novel setting. Additionally, we use historical data and naturally occurring variation to analyze heterogeneity in the *average* returns to capital. We document significant heterogeneity in returns across the $B \approx 500$ organizational divisions, both ex-ante and ex-post, with some divisions persistently underperforming expectations.

1.2.1 Measurement

Data. We construct a historic dataset by merging information on the applicant, the application, the loan terms, and the loan's performance. The dataset includes all unsecured personal installment loans issued by the bank from 2015 to 2020. Applicant data contains credit scores, while loan terms include the interest rate R, loan maturity T, face value FV, and cost of funds c. Loan performance is measured through the cost of risk q, defined as the present value of missed payments discounted at the cost of funds c, using data on the timing of default, exposure at default, recovery, and the loss-given-default.

Loan-level Returns. To measure the return to capital for issued loans, we calculate the net present value (NPV) per dollar lent, and express this as a fraction of FV. First, we compute the net present value *in the absence of default*, by discounting expected payments at the bank's cost of funds c. Next, we subtract the cost of risk q—the present value of missed payments.[8] We denote this measure π, where:

$$\pi = NPV - q$$

We estimate the *expected* π as the bank does, using publicly available credit scores. In our environment and elsewhere, the credit score provides the only signal of quality and cost of risk for underwriting and *risk-based pricing* decisions for most intermediaries. In our analysis, we use the credit score developed by Experian.

Divisional Returns. To calculate the divisional return to capital, we average π, the net present value (NPV) per dollar lent, by division. We calculate this for both expected returns, as capital allocation decisions are made ex-ante, and realized returns, reflecting actual outcomes. For ease of interpretation, we present residuals obtained by regressing

[8] The data contains information on the calendar date the loan has first reached non-performing status; as well as recovery made in the 12 month consequent to defaults for non-performing loans. We use this information to calculate the cost of risk, i.e., the present value of missed payments. First, we calculate a base PV for loans that default and those that did not. For loans that *did not* default, we calculate a base PV the usual way. For loans that *did* default, we calculate a base PV assuming uniform recovery over 12 months subsequent to the non-performing date. Second, we make adjustments to the base PV for late payments.

Table 3: Divisional Returns

Variable		unit	s.d.	iqr	p10	p50	p90
Single year (2018)							
NPV Expected	π_E	% of FV	0.50	0.49	-0.21	0.21	0.70
NPV Realized	π_R	% of FV	2.02	2.50	-3.21	-0.63	1.91
Long-run (2015-2020)							
NPV Expected	π_E	% of FV	0.31	0.36	-0.29	0.04	0.36
NPV Realized	π_R	% of FV	1.38	1.54	-1.37	0.14	1.67

Note. Based on the universe of loans issued between 2015 to 2020 across the $B \approx 500$ organizational divisions. The average return to capital over this period is $\approx 1.5\%$ of FV.

π on calendar date fixed effects, which display values relative to the mean in a given year.

Table 3 displays summary statistics for a single year and the long-run (five years), the latter to show persistence. Figure 1 displays the scatterplot of the joint distribution of expected (x-axis) and realized (y-axis) divisional performance. The figure on the left focuses on a single year and the figure on the right focuses on the long-run (five years), again to show persistence.

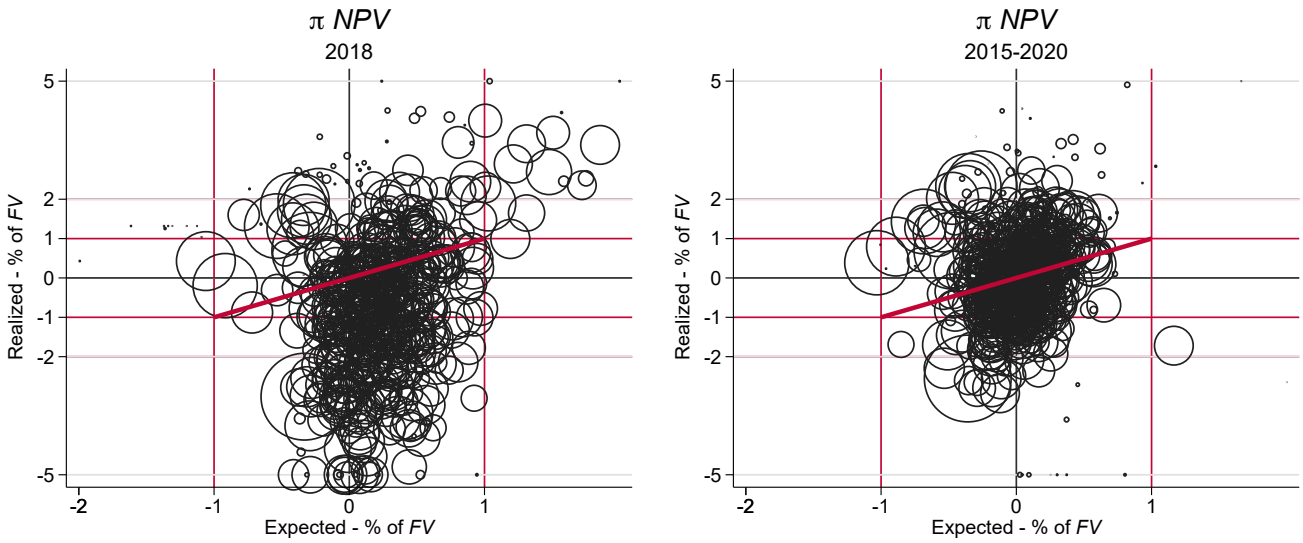
1.2.2 Heterogeneity

First, we document that returns across organizational divisions are heterogeneous. Loans originating from the worst-performing divisions have substantially lower returns compared to those from the best-performing divisions, both in expected and realized returns, and this holds true for a single year as well as over the long run.

Focusing on a single year, 2018, the division-level NPV per dollar lent at the 10th and 90th percentiles differ by 0.91% per dollar lent. If capital were allocated to equalize expected *average* products in a given year, we would observe no ex-ante heterogeneity—expectations would degenerate into a single value, resulting in a vertical line in Figure 1. Realized returns to capital are more dispersed, differing by more than 5.12% per dollar lent at the 10th and 90th percentiles.

Focusing on the long run, 2015 to 2020, allows us to average out division and year-specific conditions and business cycles trends, and reveals that these patterns are highly persistent. The average return to capital over this period is $\approx 1.5\%$ of FV. Over the long run, the division-level NPV per dollar lent ranges by more than 3% per dollar lent at the 10th and 90th percentiles—nearly twice as much as the average return on capital. On

Figure 1: Returns to Heterogeneity



Note. Scatterplots for the division level expected and realized returns to capital. Returns winsorized at 5% of FV. Each circle represents a division, with the size of the circle corresponding to its lending volume. The figure on the left focuses on a single year (2018). The figure on the right focuses on the long-run averages over a period of five years (2015-2020) in order to show persistence.

average, the realized differences in NPV between two random divisions are 1.42% of FV. A division in the bottom quartile of realized returns has a 47% chance of remaining in the bottom quartile after a year, and a 33% chance of remaining in the bottom quartile after three years.

Second, we document that some divisions *persistently underperform* relative to expectations. This can be seen in the right panel of Figure 1, which visualizes the joint heterogeneity in expected and realized performance by division. If expected performance matched realized performance, divisions would align along the 45-degree line (solid red). Otherwise, there exists a *bias* (expected-realized difference) in long-run performance. This suggests potential *sand-in-the-gears*, such as unobservable characteristics or organizational frictions, leading to realized performance dispersion beyond the expected (i.e., not captured by the credit score).[9]

This long-run heterogeneity in realized *average* returns *invites the interpretation* that the bank may be leaving money on the table by inefficiently allocating capital, with too much capital flowing to low-return, and too little capital flowing to high-return

[9]Divisions outperforming expectations (below the 45-degree line) may be benefiting from superior information (e.g., soft or local knowledge). The divisions with a low and lower-than-expected return to capital also tend to be more cyclical.

divisions. However, *average* returns do not reveal much about the *marginal* (next best) returns for each division. Next, we provide a conceptual framework to characterize these marginal returns.

2 Conceptual Framework

We now describe our approach to link *identified moments* that can be estimated via a pricing experiment to *tests of efficiency* and profit losses. Specifically, we develop a wedge accounting framework that can be applied to a wide range of questions and settings.

Consider a bank lending across B divisions. These divisions share an internal capital market and access capital at the same cost of funds, c. We do not impose a quantity constraint on the aggregate supply of capital. Our results generalize to settings where the bank is constrained by deposits and unable to increase total lending.

Let R represent the interest rate and $q(R)$ the cost of risk (i.e., the present value of missed payments). The profit for issued loans, expressed as the *NPV per dollar lent*, is given by:

$$\pi(R) = R - c - q(R)$$

The measure π, the NPV per dollar lent, 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 π if an increase in take-up, $\Theta(R)$, compensates for the reduction in R. Hence, the bank's objective is to maximize the *expected NPV per application*, defined as:

$$\Pi(R) = \Theta(R) \cdot \pi(R)$$

2.1 Optimization

Consider a division receiving a loan application and choosing a price R to maximize profits Π. While the experiment and our framework focus on pricing, they are inherently tied to capital allocation, as the optimal pricing problem is equivalent to the capital allocation problem. Deploying more capital—without altering origination quality—requires lowering the price because the demand curve is strictly downward-sloping, making price reductions and increased capital allocation equivalent.[10]

[10] Thus, the pricing experiment directly studies capital allocation by observing how deployment changes in response to pricing shocks. Pricing and capital allocation are two sides of the same coin. The equivalent optimal capital allocation problem is given by:

$$\max_K K\left(R(K) - c - q(K)\right),$$

The expected demand is $\Theta(R)$, and the cost of risk is $q(R)$. The bank's objective function and the corresponding first-order (optimality) condition are given by:

$$\max_R \Theta(R)\left(R - c - q(R)\right).$$

$$\Theta'\left(R - c - q\right) + \Theta\left(1 - q'\right) = 0 \tag{FOC}$$

In the first-order condition, the first term represents the *marginal* effect: lowering prices increases take-up, inducing an increase in capital allocation. The second term represents the *inframarginal* effect: lowering prices affects the returns on loans already taken up through changes in prices and the cost of risk. Raising prices increases profits unless the associated selection—the increase in the cost of risk—is too severe. Optimality is achieved by balancing these two effects.

2.2 Wedge

Suppose the econometrician observes an instrumental variable $\mathbb{Z}$ that is correlated with price but uncorrelated with demand Θ or costs q (and c). The optimality condition implies that if the bank was optimizing around R^* given demand and cost sensitivities, and if we use $\mathbb{Z}$ to perturb R exogenously by ΔR (as we will in our randomized experiment), it must hold that:

$$(R - c - q)\Delta\Theta + (\Delta R - \Delta q)\Theta = 0$$

Let $\sigma \equiv -\Theta' \frac{R}{\Theta}$ and $\delta \equiv q' \frac{R}{q}$ denote the elasticities of demand and costs respectively. If we use the perturbation to estimate slopes Θ' and q' and elasticities σ and δ, optimality implies:

$$\mu = \frac{\sigma - 1}{\sigma} R - \frac{\sigma - \delta}{\sigma} q - c = 0 \tag{Wedge}$$

The wedge μ is an observed gap, a distortion in the pricing first-order condition. It represents the treatment effect of a change in the price on profit, capturing both the marginal and inframarginal effects, and can be identified using exogenous variation

where $R(K)$ is the inverse demand function, that determines the maximum R at which a quantity K can be achieved. The corresponding first-order (optimality) condition given by:

$$(R - c - q) + K\left(R' - q'\right)$$

as long as Θ is invertible and differentiable. We contend with the potential of credit rationing.

in the price. To draw a parallel, in the misallocation literature, the wedge is a distortion in the capital allocation first-order condition, representing the treatment effect of a change in the input on the output, and can be identified using exogenous variation in the inputs.

As a result of these wedges, the bank prices (and allocates capital) differently than what would be optimal. The existence of the wedge is indicative of inefficiencies: efficiency requires that the wedge be equal to zero—i.e., the chosen price satisfies the first-order condition. To the extent that $\mu \neq 0$, the bank can increase profits by adjusting interest rates. A positive wedge indicates that the bank should increase the price which would imply a decrease in capital allocation, provided the demand curve is downward-sloping. Conversely, a negative wedge indicates that the bank should decrease the price which would imply an increase in capital allocation. Hence, the only impediment to efficient resource allocation is the wedge.

If demand is perfectly elastic ($\Theta' \to \infty$), there is no scope for optimizing the interest rate. In this case the wedge equals the (average) return to capital—$\mu = \pi = R - c - q$. In the case of no selection ($q' = 0$) the wedge simplifies to $\mu = \frac{\sigma-1}{\sigma} R - c - q$.

2.3 Loss

To calculate how much money is left on the table, our approach provides an analytical formula to express profit losses as a function of the cross-sectional variance of wedges across divisions. To derive the impact of wedges on profits, we calculate profit loss $\mathcal{L} = \Pi(R) - \Pi(R^*)$ as a second-order Taylor expansion of Π on R around the optimal interest rate R^*. This exercise is similar in spirit to the second-order approximation

of misallocation costs in general equilibrium as in Baqaee and Farhi (2020).[11]

$$\begin{aligned}\mathcal{L} &= \Pi(R) - \Pi(R^*) \\ &= -\frac{\mu^2}{2}\xi \\ &= -\frac{\mu^2}{2}\frac{\Theta'^2}{\Theta q'' - 2(1-q')\Theta' - (R-c-q)\Theta''}\end{aligned} \tag{Loss}$$

The profit loss yields a function of the wedge μ as well as first and second-order slopes of demand and costs: Θ', Θ'', q', q''. Notably, the availability of variation that both increases and decreases prices is crucial to estimate Θ'' and q'' without parametric assumptions.

If the bank was optimizing considering demand and cost sensitivities, then $\mu = 0$, and there is no loss. Otherwise, profit losses are quadratic in wedges: small distortions around the optimum entail approximately zero loss, but efficiency losses increase with wedges. The loss is zero if demand is perfectly elastic ($\Theta' = 0$). Under linear demand and cost curves (i.e., $\Theta'' = q'' = 0$), the expression simplifies to $\xi = \frac{\Theta'}{-2(1-q')}$.

Let $\mathcal{L}_b$ denote the loss at the branch level. The loss at the bank level is given by:

$$\mathcal{L}_{\text{Bank}} = \sum_b \mathcal{L}(\mu_b) = -\sum_b \frac{\mu_b^2}{2}\xi_b \tag{Loss, Total}$$

2.4 Efficiency: Unconstrained vs. Constrained

We say an *allocative inefficiency* exists if the bank (given its information) could use pricing adjustments to increase overall profits. This framework allows us to distinguish between two notions of efficiency.

[11] Quantifying the total profit loss, up to the second order, involves summing a series of triangles, where the height corresponds to the change in the wedge, and the base reflects the change in loans issued as a result of this wedge, which can be expressed as a function of the demand and cost elasticities.

$$\begin{aligned}\Pi(R+\varepsilon) &= (R+\varepsilon-c-q(R+\varepsilon))\,\Theta(R+\varepsilon) \\ &= \left(R+\varepsilon-c-\left(q+q'\varepsilon+\frac{q''}{2}\varepsilon^2\right)\right)\left(\Theta+\Theta'\varepsilon+\frac{\Theta''}{2}\varepsilon^2\right) \\ &= \left((R-c-q)+(1-q')\varepsilon-\frac{q''}{2}\varepsilon^2\right)\left(\Theta+\Theta'\varepsilon+\frac{\Theta''}{2}\varepsilon^2\right) \\ &= (R-c-q)\Theta+\left[(R-c-q)\Theta'+\Theta(1-q')\right]\varepsilon+\left[(R-c-q)\frac{\Theta''}{2}+(1-q')\Theta'-\frac{q''}{2}\Theta\right]\varepsilon^2 \\ &= \Pi(R)+\frac{\varepsilon^2}{2}\left[(R-c-q)\Theta''+2(1-q')\Theta'-\Theta q''\right]\end{aligned}$$

Unconstrained efficiency tests whether divisional decisions are optimal in isolation, requiring that wedges are equal to zero for each division—i.e., each division optimizes its first-order condition. Formally, testing for unconstrained efficiency involves the null hypothesis that the wedge is zero for all divisions:

$$H_0 : \mu_b = 0 \quad \text{for all } b \qquad \text{(Test, Unconstrained)}$$

—i.e., degeneracy at the value of zero for all divisions. If $\mu_b \neq 0$, the bank can increase profits by adjusting interest rates. A positive wedge indicates that the bank should raise the price, while a negative wedge suggests it should lower the price.

Constrained efficiency, on the other hand, evaluates the potential gains from reallocating capital across divisions, recognizing that divisions operate within a shared internal capital market and cost of funding c. Hence, efficiency requires that the marginal return to capital is equalized across divisions.This is a *factor-augmenting* reallocation across divisions that holds lending constant but generates more profits.

Constrained efficiency requires that wedges are equal to a common value across divisions, though this common value is not necessarily zero. If one division has a higher wedge (indicating higher profitability from raising prices), while another has a lower wedge (indicating profitability from lowering prices), then reallocating capital can increase overall profits. A common wedge value ensures that all divisions are operating under the same pricing distortions, eliminating inefficiencies that arise from variation in pricing distortions.

Formally, testing for constrained efficiency involves the null hypothesis that the variance of wedges across divisions is zero $var(\mu_b) = 0$, equivalently, that the total profit loss is zero:

$$H_0 : \mathcal{L}_b = 0. \qquad \text{(Test, Constrained)}$$

To the extent that μ_b is not constant across divisions, the bank can improve overall profitability by adjusting interest rates to reduce or eliminate the wedges.

3 Pricing Experiment

We now use a randomized experiment to identify the wedge ingredients. Conceptually, the goal is to induce more or less capital allocation at a division by randomizing prices, either lower or higher than the baseline, and measure the profits when loans that do not get issued in naturally occurring data are issued. Econometrically, this approach requires an instrument $\mathbb{Z}$ that is correlated with the supply (e.g., price) of allocated capital but uncorrelated with demand or costs.

The experiment randomizes prices, allowing us to estimate the treatment effects of price changes on demand and costs. We use the randomized assignment to pricing groups as an instrument that overcomes the endogeneity of prices (and hence the allocation of capital). This section presents descriptive evidence on average first-stage, intent-to-treat, and treatment effects. In the next section, we then use the randomized prices to test for efficiency and estimate the implied profit loss.

Design. In the experiment, consumers were randomly allocated to three groups. $\mathbb{Z} = \{-1, 0, 1\}$ corresponds to the group assigned to the borrower. These groups were offered different interest rates R. The randomization was based on the digits of the application number. There is no explicit participation choice; therefore, applicants are not aware they are participants in a controlled trial.

The baseline group (80%) experienced no change in price compared with the status quo. The low-price (L) group, consisting of 10% of borrowers, was offered a price reduction. The high-price (H) group, also consisting of 10% of borrowers, 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.

The experimental procedure provides two randomized instruments, $\mathbb{Z}^L$ and $\mathbb{Z}^H$, for econometric evaluation. This variation allows us to observe behavior in response to both price increases and decreases. This is a key ingredient in our profit loss calculation and enables the estimation of second derivatives, which are central to the second-order approximation.

Summary Statistics. Table 4 displays summary statistics. Panel A focuses on application characteristics, while Panel B focuses on outcomes. The average applicant is offered an interest rate of 18.8%, with an average term of 29 months and a requested face value of 29,611 TRY. The maximum term is three years, with about two-thirds of all loans having a term of three years.

Our key outcome variables are demand $\Theta \in \{0, 1\}$—i.e., take-up—and the cost of risk $q \in [0, 1]$—i.e., the present value (PV) of missed payments (assuming that the bank's cost of funds is R^c) normalized by the face value of the loan.

A total of 17,096 applications (34.2%) were taken up by applicants and issued as loans. This subset allows us to assess loan performance (i.e., the present value of payments) and estimate cost sensitivities. Among loans taken up, 3.7% fall into 90+ (nonperforming) status, resulting in an average cost of risk of 1.72% of the face value. This yields a realized NPV of 1.73% of the face value for issued loans and a NPV per application of 0.59 % of FV.

Table 4: Experiment: Summary Statistics

Variable		unit	N	mean	s.d.	p10	p50	p90
Panel A: Application								
Metro			50,042	0.54	0.50	0	0	1
Loans		Count	50,042	1.6	1.7	0	1	4
Score			50,042	1,407	237	1,077	1,415	1,728
Amount	FV	TRY	50,042	29,611	28,914	6,500	24,000	50,000
Term	T	Months	50,042	29.0	9.4	12	36	36
Cost of Funds	c	% APR	50,042	15.7	0.7	14.9	16.2	16.2
Spread	R^s	% APR	50,042	3.2	0.9	2.3	2.9	4.2
Price	R	% APR	50,042	18.8	1.1	17.4	18.7	20.4
$\mathbb{Z}^L$			50,042	0.10	0.30	0	0	0
$\mathbb{Z}^H$			50,042	0.10	0.30	0	0	0
Panel B: Outcomes								
Take-up	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
NPV	π	% of FV	17,096	1.73	9.81	1.29	3.53	5.40
NPV per App.	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
Other costs	ζ	% of FV	17,096	0.02	0.07	0	0	0.07

Table 5: Experiment: Covariate Balance

	Metro	**Loans** Count	**Score**	T Months	R^c % APR
$\mathbb{Z}^L$ Low	0.53 (0.007)	1.62 (0.02)	1428 (3.4)	28.8 (0.13)	15.67 (0.01)
Baseline	0.54 (0.003)	1.63 (0.08)	1407 (1.2)	29.1 (0.05)	15.67 (0.003)
$\mathbb{Z}^H$ High	0.52 (0.007)	1.64 (0.02)	1391 (3.4)	28.9 (0.14)	15.71 (0.01)
N	50,042	50,042	50,042	50,042	50,042

Balance. Table 5 displays the covariate balance table. The averages for the high, low, and baseline groups are indistinguishable.

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

used only to the random assignment of a participant to the control versus treatment group, $\mathbb{Z}_i$.[12]

Table 6: Experiment: Instrumental Variables Estimation Framework

	Z Instrument	**X** RHS	**Y** LHS
First-stage	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	$\mathbb{Z}^L$ and $\mathbb{Z}^H$	R
Intent-to-treat	$\mathbb{Z}^L$ 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

Estimation Framework. Table 6 summarizes the estimation framework used to analyze the randomized controlled trial. First, we employ first-stage and intent-to-treat (ITT) specifications to analyze the effects of the experiment on prices, demand, and costs.

In the first stage, Y is R, the interest rate. In the ITT specifications, Y is either a take-up dummy (demand) or the present value of missed payments normalized by the face value (costs). The instruments $\mathbb{Z}$ are orthogonal to measurement error, omitted variables, and the residuals.

To estimate a treatment effect within a two-stage least-squares (2SLS) framework, we replace the right-hand-side variable with the price R, using $\mathbb{Z}^L$ and $\mathbb{Z}^H$ as instruments. 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 ITT and first-stage effects:

$$\beta^{\text{Wald}} = \frac{\beta^{\text{ITT}}}{\beta^{\text{FS}}}$$

Average Effect Estimates. Table 7 displays the first-stage (FS) and intent-to-treat (ITT) specifications and compares the average change in price and outcomes (demand or cost) between the treatment group and the control group using ordinary least squares (OLS).

Focusing on the first-stage estimates in the first column of Table 7, the baseline

[12] To ensure the experiment only affects prices, we investigate whether price changes trigger a revision in the final loan amountand term. We do not find such an inclination. About 6% of the applicants change their term, and 4% change their loan amount by more than 10%. However, these changes are orthogonal to the experimental assignment.

Table 7: Experiment: First-stage and Intent-to-treat Estimates

	First-stage		Intent-to-treat		Wald Treatment Effect per %p APR	
	R **Price** % APR	R^s **Spread** % APR	Θ **Take-up** %	q **Cost-of-risk** % of FV	Θ' **Take-up** %	q' **Cost-of-risk** % of FV
$\mathbb{Z}^L$ Low	-1.32	-1.32	13.2	-0.72	-10	0.55
	(0.01)	(0.01)	(0.7)	(0.22)		
Constant	18.89	3.22	34.5	1.80		
	(0.005)	(0.003)	(0.2)	(0.08)		
$\mathbb{Z}^H$ High	0.92	0.88	-16.6	0.38	-18	0.41
	(0.01)	(0.01)	(0.7)	(0.34)		
N	50,042	50,042	50,042	17,096		
F	7,286	12,602				

group receives an interest rate of 18.9% APR, which is 1.32%p lower for group $\mathbb{Z}^L$ and 0.92%p higher for group $\mathbb{Z}^H$. The F-statistic for the first stage is 7,286. As the cost of funds c is balanced across the groups, the difference in the interest rate comes entirely from the spread.

The intent-to-treat estimates reported in the third and fourth columns of Table 7 show that lowering prices increases take-up Θ, and vice versa. The take-up level for the baseline group is 34.5%, which is 13.2%p (38%) higher for group $\mathbb{Z}^L$ and 16.6%p (48%) lower for group $\mathbb{Z}^H$. In response to a 1%p APR decrease in prices (from a base of 18.9% APR), demand increases by 10%p (29% off a base of 34.5%). Conversely, in response to a 1%p APR increase in prices, demand decreases by 18%p (52%).

The treatment effect on demand is the marginal effect of price changes that induce more (or less) capital allocation at a division. However, price changes also generate an inframarginal effect, influencing the returns on loans already taken up through changes in prices and the cost of risk. Notably, lower prices attract borrowers with better risk profiles, while higher prices deter safer borrowers, leaving a pool of riskier loans. The cost of risk q (present value of missed payments)—which, for the baseline group, is 1.80% of FV—is 0.72% of FV (40%) lower for group $\mathbb{Z}^L$ and 0.38% of FV (21%) higher for group $\mathbb{Z}^H$. Hence, in response to a 1%p APR decrease in prices, cost of risk decreases by 0.55 of FV. Conversely, in response to a 1%p APR increase in prices, demand decreases by 0.41 of FV.

4 Results

We now operationalize the apparatus in Section 2 to test the efficiency of capital allocation and quantify the resulting profit losses. First, we calculate wedges based on experimentally identified treatment effects of price changes on demand and costs. Then, we estimate the bank-level profit loss implied by the distribution of these treatment effects. Below, we discuss how to obtain estimates of these objects.

The experiment involves 50,042 applications across $B \approx 500$ branches. To estimate heterogeneous demand and costs, we aggregate branches into groups. Grouping branches ensures a sufficient number of observations in each group, avoiding statistical issues such as zero costs (no defaults) and imprecisely estimated cost sensitivities. We then calculate wedges for each group, as well as the dispersion across these groups. As explained in Section 2, our tests are based on the properties of the distribution of wedges rather than individual wedges.

We group branches based on the *bias* variable (the long-run difference between expected and realized performance) discussed in Section 1.2. Thus, when reallocating, we do so between branches that underperform relative to expectations and those that overperform. This approach ensures that reallocation is orthogonal to observable characteristics. For instance, across bins, variation in other features such as credit score and contract term is minimal, with interquartile ranges of only 0.6% and 1% of their respective averages. Consequently, we reallocate loans across divisions rather than based on borrower characteristics.

Recall that the wedge μ is a function of the interest rate R, cost of funds c, demand (take-up) Θ, cost of risk q, and the identified first and second derivatives of demand and costs. R and c are observed ex-ante at the loan level and are unchanged ex-post. Θ and q are observed ex-post and can also be estimated ex-ante. We estimate Θ' and q' using the experimental design and instrumental variables approach described in Section 3. We calculate Θ'' and q'' from the first derivatives using finite differences. Hence, our approach does not require assumptions about the functional form of demand or costs.

In the misallocation literature, capital is deemed misallocated if a planner facing the *same constraints and the same information* could reallocate inputs across firms and obtain a higher output. Similarly, our definition of allocative efficiency is given the same constraints and information that the bank faces. Capital allocation decisions are made ex-ante, based on the information available at the time. Realizations include a noise or luck component and the heterogeneity in realized returns may reflect this. Hence, the ex-ante dispersion in the wedges is not equal to the ex-post dispersion. We report expected values as our main estimates, as is done in this literature.

When calculating expected values, we consider two scenarios. First, we calculate these values from the bank's perspective. For example, the bank does not have a model for demand. Second, we calculate expected values from the econometrician's perspective. In this case, we condition on variables observable to the econometrician. These two approaches answers two 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?

Finally, to calculate profit losses for the bank's total lending, we apply lending weights (the share of lending as a fraction of overall lending in the year prior) and estimate the lending-weighted mean of the wedges and loss.

4.1 Unconstrained Efficiency

First, we test for *unconstrained* efficiency. Recall that the wedge is the treatment effect of a price change on profit, capturing both marginal and inframarginal effects, given by Equation Wedge. The test of unconstrained efficiency is that for any division, the wedge is equal to zero, or $H_0 : \mu_b = 0$—*degeneracy at the value of zero for all divisions*.

Table 8 display the results. The average treatment effect of a supply-driven price decrease on profit per application is different than zero. We estimate the mean of wedges to be 0.87%, with a standard error of 0.49% and a p value of 0.09. Hence the test for unconstrained efficiency does not reject. As the estimated demand sensitivity $\sigma > 1$ and hence $\frac{\sigma-1}{\sigma} > 0$ for every group, a positive wedge indicates that the branch is *underpriced*. As the demand curve is downward sloping, this implies capital is over-allocated to the division. The average wedge exceeds zero, and hence, on average, there are gains from raising prices.

The wedges take both positive and negative values. Approximately half (48%) of the divisions exhibit a negative wedge, suggesting the division is overpricing—the price should be lowered, which would increase the allocated capital. For the remaining divisions, a positive wedge is indicating the division is underpricing, and hence a need to decrease the allocated allocated capital. The test of unconstrained efficiency—that for any division, the wedge is not equal to zero—does not tell how large are the profit losses. We next proceed with quantifying the loss.

4.2 Constrained Efficiency and Profit Loss

Testing for *constrained* efficiency $H_0 : \mathcal{L}_b = 0$ is analogous but takes into account that the divisions are connected through the internal capital market and share a common cost of funds c. Hence, a constrained efficient allocation would feature *wedges*

Table 8: Experiment: Results

	μ **Wedge**		$\mathcal{L}$ **Loss**		
	Estimate	p	Estimate	p	%
Panel A: Main Estimates					
Expectation	0.87	0.09	-0.109	0.001	17.5%
	(0.49)		(0.029)		
Expectation of Bank	0.66	0.18	-0.085	0.0004	
	(0.47)		(0.020)		
Realization	0.72	0.16	-0.062	0.19	
	(0.49)		(0.046)		
Panel B: Main Estimates—Cross-sell and Issuance Costs					
Expectation	1.09	0.04	-0.115	0.0008	18.5 %
	(0.49)		(0.029)		
Realization	0.89	0.07	-0.079	0.06	
	(0.46)		(0.039)		
Panel C: Extensions					
Across Scores	-0.22	0.61	-0.046	0.087	7.4%
	(0.39)		(0.020)		
Across Divisions Within Scores	0.37	0.48	-0.098	0.008	15.8%
	(0.53)		(0.034)		
Panel D: Main Estimates—Robustness					
Winsorize at 5%	0.91	0.10	-0.111	0.001	17.8 %
	(0.53)		(0.030)		
50 groups	0.91	0.009	-0.116	0.000002	18.6 %
	(0.33)		(0.021)		
Excluding Online Branch	0.97	0.08	-0.105	0.003	16.9%
	(0.52)		(0.032)		

equal to a common value across divisions, though this common value is not necessarily zero. Constrained efficiency evaluates the potential benefits of reallocation, and hence

concerns the heterogeneity in the marginal return—the dispersion in wedges. If one division has a higher wedge (indicating higher profitability from raising prices), while another has a lower wedge (indicating profitability from lowering prices), then reallocating capital can increase overall profits. A common wedge value ensures that all divisions are operating under the same pricing distortions, eliminating inefficiencies that arise from variation in pricing distortions.

The profit loss $\mathcal{L}$ is a function of the wedge μ and wedge ingredients Θ' and q', as well as the second order identified slopes of demand and costs: q'' and Θ''. We calculate Θ'' and q'' from the first derivative using finite differences. Hence, an experiment that both hikes and cuts prices is crucial for estimation without assumptions on the functional form of demand Θ and costs q.

We decisively reject the null of constrained efficiency, with of $p = 0.001$. Hence, there are benefits to reallocating capital across divisions until wedges are equal.

The constrained efficiency calculation allows us to go beyond testing to quantify the profit losses in economic terms. As wedges are dispersed across divisions implies gains from reallocating capital to higher marginal return divisions, holding lending constant. This is a *factor-augmenting* reallocation across divisions that holds lending constant but generates more profits.

We find the gain from potential reallocation through equalizing wedges to be substantial—0.109 % of FV. Compared to the average expected and realized profit per application Π of 0.623% per dollar lent this implies a loss of 17.5%.[13]

4.3 Robustness and Extensions

We provide several robustness checks and extensions for our wedge and loss calculations. These results are reported in Panels B and C of Table 8.

Cross-sell and Issuance Costs. Our analysis so far ignores cross-sell and issuance costs. First, 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 borrowers who later purchase other products. This cross-subsidization could allow banks to optimize overall profitability, even if individual decisions and divisions appear inefficient in isolation.

[13] As we later discuss, these numbers are broadly compatible with profit losses within companies documented elsewhere. Most notably, DellaVigna and Gentzkow (2019) show that most U.S. food, drugstore, and mass-merchandise chains price uniformly across divisions, resulting in a median store profit loss of 1.79% of *revenue*, which translates to 17.9% of profits at a 10% operating margin and 35% of profits at a 5% operating margin. Giroud et al. (2022) uses plant-level data from the U.S. Census Bureau, and show that in a counterfactual where all multi-unit companies are assigned the highest management score, aggregate productivity increases by 13% due to improvements in the allocation of internal resources.

Second, 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 marginal costs associated with originating loans, such as application, origination, and underwriting costs at the branch and headquarters. 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 other costs by ζ. Then the NPV per dollar lent 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 (marginal) costs in the same way the bank accounts for them. These marginal costs (about 0.02% of FV) are negligible in magnitude compared with the profit per dollar lent. Moreover, they do not vary across divisions and, therefore, do not affect constrained efficiency (i.e., profit loss). The third row of Panel C reports these results. We find results similar to our main estimates.

Inefficiencies Across Credit Scores. Our analysis reveals inefficiencies in pricing—and, consequently, capital allocation—*across divisions*. A natural question arises: is the bank pricing efficiently within the groups it uses as a basis for pricing, such as credit score groups? This analysis addresses concerns about whether the bank can optimize effectively, even when allowed to condition on characteristics it is capable of optimizing.

Our bank prices loans according to standard a risk-based cost-plus method. This approach adds a fixed markup to total costs (i.e., cost of funds and cost of risk) to determine the spread between the loan's interest rate and the cost of funds, targeting a rate of return. Hence, the bank charges higher prices that include a risk premium for borrowers with higher costs of risk to compensate for the probability and loss given default. However, such risk-based pricing fails to set optimal prices as it overlooks key variations in borrowers' price sensitivities.

To provide additional context on profit losses across divisions, we use our framework to evaluate inefficiencies within the bank's existing risk-based pricing system and test whether pricing is efficient across credit score groups. The results of this analysis

are shown in the first row of Panel C. Recall that the bank assigns all borrowers to one of five spreads based on credit score. We group loans by division *within* these risk-based pricing categories. We we perform the same analysis for groups based on credit score, we decisively fail to reject the null hypothesis of unconstrained efficiency ($p = 0.96$). Additionally, while we estimate a profit loss of 7%, we again fail to reject the null hypothesis of constrained efficiency ($p = 0.093$). This profit loss is primarily due to the bank's risk-based pricing overpricing borrowers with the lowest credit scores, potentially due to a failure to account for how cost of risk changes as price changes (e.g., adverse selection) within this group.

As a second check, we perform a more nuanced version of the above analysis, reallocating loans across divisions but within credit score groups (thus isolating risk). We group loans by division *within* these five risk-based pricing categories and assess the benefits of reallocating across divisions but within these groupings. The results of this analysis are reported in the second row of Panel C. Similar to our main results, we reject the null hypothesis of unconstrained efficiency ($p = 0.48$) but fail to reject constrained efficiency ($p = 0.008$), and we estimate a profit loss of 15.8%. Hence, capital allocation remains inefficient across divisions, even after reallocating within credit scores.

Other. We also provide a number of robustness checks. First, the quadratic approximation to the loss is sensitive to outliers. In our baseline estimation, we winsorize the sample at 10% of the wedge. As a robustness check, we also report calculations where we winsorize at 5%. Second, instead of grouping divisions into 20 groups in our baseline estimation, we use 50 groups. Third, we exclude the online branch. These results are presented in Panel B of Table 8. We find broadly consistent results, with profit losses of 17.8%, 18.6%, and 16.9%; and p-values (for the tests of constrained efficiency) of 0.01, 0.000002, and 0.003, respectively.

Overall, we find that wedges are slightly larger than zero (inconsistent with unconstrained efficiency), however we fail to reject the null that, on average, bank's optimize their first order condition. However, we also find that wedges display substantial dispersion across divisions (inconsistent with constrained efficiency), which points to benefits from reallocation until wedges are equalized.

5 Discussion

5.1 Why is the Capital Allocation Inefficient?

It may be surprising that a single bank fails to allocate capital efficiently *within*. After all, a single internal capital market does not face the commonly invoked financial

frictions and by virtue of its control rights, can actively reallocate capital across divisions in ways that an *external* market cannot. So, why is the pricing—and, consequently, the allocation of capital—inefficient?

Comprehensively identifying all the root causes contributing to profit losses is beyond the scope of this article. Potential *sand-in-the-gears* factors run the gamut: incentives, authority, information, and competition may lead to both heterogeneous returns across divisions and persistent underperformance of some divisions.[14] However, the price mechanism could still be used to achieve allocative efficiency by *reallocating* capital on the margin—redirecting capital away from low-return divisions to high-return divisions. After all, recall from the framework in Section 2 that efficiency only requires prices to be adjusted to eliminate wedges. Why, then, are prices not adjusted to achieve efficiency?

We argue that banks fail to achieve efficiency because they engage in risk-based pricing that is *uniform* across divisions. In other words, banks set a single price across all divisions, failing to account for local factors, divisional unobservables, and organizational frictions that cause some divisions to underperform. 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, 2021), and (Begenau and Stafford, 2023), and mortgages securitized by government-sponsored enterprises (Hurst et al., 2016).

To illustrate this further, consider the variation across divisions in the price R compared to the demand Θ, demand slope Θ', cost of risk q, and cost of risk slope q' (i.e., selection). For the price, the interquartile range across divisions is merely 0.2% of the average value. In contrast, for Θ, Θ', q, and q', the interquartile ranges are strikingly 6%, 25%, 31%, and 138% of their respective averages. Thus, the variation in demand is *30 times*, demand slope and cost of risk is approximately *100 times*, and the variation in cost of risk slope is more than *500 times*, that of the price.

Heterogeneity in divisional performance may be driven by unobservable characteristics: the same credit score may imply a higher default rate for one division compared to another due to differences in the customer base across geographies. Alternatively, differences in realized performance could stem from operational or organizational frictions, such as branch managers soliciting unobservably low-return applications. Our analysis does not aim to disentangle these competing explanations.

However, prices should serve as a corrective tool to offset these unobservable or

[14] Division managers—who are incentivized based on lending quantity rather than profitability—may misuse their discretion in loan approvals, fail to fully incorporate local or soft information, or approve loans at prices that are too low. Conclusively identifying these frictions is beyond the scope of this article.

organizational frictions causing divisional underperformance. Efficiency only requires prices to adjust to eliminate wedges. Nevertheless, uniform pricing practices lead to inefficient capital allocation by preventing banks to condition on geography and hence using prices from being used as a tool for reallocation. To achieve optimal capital allocation across divisions, banks should adopt more flexible pricing strategies that respond to organizational frictions, as well as local demand and cost variations.

To reiterate, uniform pricing is *not* the root cause of inefficiency but rather an *impediment to reallocation*. The inefficiency in capital allocation persists because prices fail to adjust to facilitate capital reallocation. Thus, while our framework does not address why divisions initially have low or high returns, it explains why these inefficiencies remain uncorrected.

5.2 Why is Pricing Uniform?

These findings are puzzling. Why do banks price uniformly across divisions and fail to account for division-level factors in their pricing decisions? Are they unable (e.g., due to regulations) or unwilling (e.g., due to managerial inertia or brand image concerns) to set different prices?

Traditional menu cost models cannot explain uniform pricing across divisions in banks. Banks operate in a highly digital pricing environment where interest rates on loans and deposits are updated in real-time based on borrower risk and market conditions. The technological infrastructure of banks enables automated, algorithm-driven pricing adjustments, minimizing the cost of implementing geographic or division-specific rate differentiation. Moreover, unlike retail chains, our bank does not engage in highly publicized national *billboard* pricing, where prices are publicly advertised and standardized across locations to maintain a consistent brand image. Finally, this uniformity also does not arise from securitization, as in the case of government-sponsored U.S. mortgage market.

The defining factor behind uniform pricing is the necessity to comply with legal protections against discriminatory lending. These regulations are known in the U.S. as *redlining* regulations. Redlining refers to the practice of denying individuals in certain areas equitable access to credit, through discriminatory pricing or underwriting. In the U.S., redlining originated as a practice of denying credit or charging higher prices to individuals in predominantly Black neighborhoods, often codified through racially discriminatory lending maps. While there are no direct redlining regulations in the Europe Union, various competition laws (e.g., ensuring fair access), anti-discrimination laws (e.g., prohibition of discrimination based on protected characteristics), consumer protection laws (e.g., geo-blocking regulation), and housing laws collectively ensure

equitable access and effectively regulate practices that could be considered *place-based discrimination*—exclusionary based on geographic factors. Financial institutions must comply with these laws to ensure equitable access to services. Hence, concerns about equity and discrimination can lead to uniform pricing practices.

In Appendix C and Table 10, we comprehensively discuss these laws. In the U.S., Fair Housing Act (FHA) (42 U.S.C. §§3601-3619) prohibits redlining and lending discrimination based on location. In the European Union, Payment Accounts Directive establishes a right to basic banking services and explicitly bans geographic discrimination in that context. In France, Penal Code (Article 225-1) lists "le lieu de résidence" (place of residence) among protected grounds alongside origin, sex, disability, etc. and hence criminalizes place-of-residence-based discrimination in lending. In Turkey, Law No. 6701 prohibits both public and private entities from discriminating in offering services, which includes banking and finance.[15]

Regulations designed to prevent geographic discrimination restrict financial institutions from conditioning prices on geographic information. If allowed, region-specific risk factors could improve capital allocation. However, by being unable to adjust for regional differences, institutions are forced to treat all divisions similarly, even when theory and data suggests otherwise. While these policies may promote fairness, they prevent banks from implementing optimal prices, and hence preclude the use of prices as a tool to optimize capital allocation, resulting in inefficiencies.

What is the Wedge and Loss Related To?

Although empirically distinguishing the various frictions is beyond the scope of this article, analyzing the correlates of the wedge and the loss provides some suggestive evidence. The main correlate of the profit loss is the *bias*, the expected-realized difference in long-run performance, discussed in Section 1.2. This variable captures the *sand-in-the-gears*, such as unobservable characteristics or organizational frictions, that make some divisions *persistently underperform* relative to expectations, leading to cost and return realizations beyond the expected. We find that a 1 % increase in the bias

[15] In contrast to loan pricing, financial institutions around the world generally have more flexibility in setting deposit rates across geographies. For example, in the U.S., there are no federal anti-discrimination laws for deposit pricing: unlike loans, deposit pricing is not subject to regulations such as the Equal Credit Opportunity Act (ECOA) or the Fair Housing Act (FHA). Similarly, in the E.U., deposit pricing is not covered by the Mortgage Credit Directive (MCD) or anti-discrimination provisions under broader E.U. directives like the Equal Treatment Directive. Moreover, unlike loans, deposit pricing is not subject to interest rate caps (i.e., usury laws), and the Community Reinvestment Act (CRA) does not directly regulate deposit pricing—banks may offer higher deposit rates in underserved areas to meet CRA obligations. Hence, banks *may* set different savings and certificates of deposit rates in different markets based on branch presence, deposit demand, local competition, customer acquisition strategies, and branch operating costs, offering lower rates in remote and less competitive markets (e.g., North Dakota) compared to highly competitive, urban financial hubs (e.g., New York City).

leads to an increase in divisional loss of . . . %.

In contrast, we do *not* find a correlation between the *average* returns and the loss. Hence, using naturally occurring data on average returns, in our context, is not informative about the efficiency of capital allocation.

5.3 Broader Implications

5.3.1 Redistribution

Uniform pricing, compared to optimal pricing, changes both the pricing and the allocation of capital across divisions. This has redistributive effects, generating winners and losers and impacting inequality. However, the direction of these effects may not be immediately clear. Which divisions are subsidized—receiving a lower price and more capital than they should? Who are the winners and losers?

To understand the effects on pricing and capital allocation, we compare the optimal and uniform prices (and the implied quantities). Calculating the optimal prices for each division provides a counterfactual scenario, illustrating how interest rates and capital allocation would change. Using Equation (Wedge), we can calculate the optimal price as:

$$R^* = \frac{\sigma}{\sigma - 1} c + \frac{\sigma - \delta}{\sigma - 1} q$$

where the elasticities σ an δ are identified using instrumental variables.

If the optimal price aligns with income, uniform pricing will lower prices for high-income divisions. Conversely, if it aligns with (unobservable) risk, uniform pricing will lower prices for (unobservably) high-risk divisions.

First, the optimal price is negatively correlated with income: lower-income borrowers would have faced higher prices under non-uniform pricing but instead pay lower prices under uniform pricing. We find that a 1% increase in income is associated with approximately a 0.35% decrease in the optimal price. This implies that uniform pricing redistributes capital by subsidizing lower-income borrowers at the expense of higher-income borrowers, potentially reducing inequality.

Second, borrowers in divisions where returns to capital exceed expectations (e.g., persistent overperformance or a high *bias*, as discussed in Section 1.2) are overcharged under uniform pricing. Specifically, a 1% increase in bias—defined as the division-level long-run realized NPV per dollar lent above expectations—is associated with about a 0.02% *decrease* in the optimal price under non-uniform pricing. Conversely,

borrowers in underperforming divisions face higher prices under non-uniform pricing compared to uniform pricing. This redistribution mirrors the dynamics of socialistic internal capital markets, where resources are transferred from better-performing to underperforming divisions, effectively subsidizing weaker divisions at the expense of stronger ones, as in Scharfstein and Stein (2000). Uniform pricing leads to inefficient cross-subsidization, underpricing underperforming divisions.

In addition to the legal risk, brand image concerns—mechanisms by offering borrowers with similar risk different prices at different branches—may lead to negative reactions, making them think it is unfair and depressing demand in the long-run.

5.3.2 Pass-through

How local prices adjust to local economic conditions is also central to a range of economic policy questions, from understanding business cycles to providing economic relief. A second implication of uniform pricing is that it dampens how prices, and hence the allocated capital, respond to changes in local conditions. To illustrate, consider an organizational shock at the division level that increases the cost of risk by q. The optimal price R^* should then adjust upward to account for the higher risk, reallocating capital away from the affected division and toward divisions with higher returns. Under uniform pricing, such adjustments do not take place. As a result, divisions with elevated risk receive prices too low and capital too high, reducing efficiency from the bank's viewpoint. This rigidity, which prevents prices from not adjusting efficiently, also creates a basis for local government policy response.

5.4 Future Work

Our study suggests several avenues for future research. First, decomposing the sources of inefficiencies: potential *sand-in-the-gears* span incentives, authority, information, and competition, all of which may drive persistent heterogeneity in returns across divisions. Identifying the relative contribution of each factor would shed light on the root causes of capital misallocation. Second, *what is the best way to correct misallocation?* A range of levers could be pulled to allocate capital to high-return divisions—promotional campaigns, employee incentives, better incorporation of local or soft information, adjustments to managerial discretion, or modifications to underwriting thresholds. Each of these mechanisms reallocates capital through distinct marginal projects and interacts differently with inframarginal loans (e.g., relaxing a quality threshold would entail no inframarginal effects). Future work could benefit from designing experiments to compare these strategies and evaluate their effectiveness in improving capital allocation.

6 Conclusion

In this study, we studied a fundamental question—whether capital is allocated within a financial institution to equalize marginal returns across divisions. To do so, we developed a novel methodology that uses pricing variation to test the efficiency of capital allocation and quantify the resulting losses. Our approach avoided restrictive assumptions about functional forms, the extent of optimizing behavior, or the underlying sources of inefficiencies.

We applied these methods to study bank internal capital markets that allocate deposits across branches at a shared cost of funds. We used a high-stakes, large-scale randomized experiment that facilitates capital reallocation through pricing shocks. These pricing shocks, by inducing more (or less) capital allocation, enabled us to identify the *marginal* returns to capital and apply our methodology. We traced the response of capital deployed and profits generated in response to these random shocks by merging administrative data on take-up and realized performance. Our tests of unconstrained efficiency indicated detectable deviations from efficiency. The tests of *constrained efficiency* gave a profit loss of 17%.

We argued that banks fail to efficiently allocate capital across divisions because they engage in risk-based pricing that is *uniform* across divisions, without accounting for unobservable characteristics and organizational frictions that cause some divisions to persistently underperform expectations. We find that the variation in realized demand sensitivity and cost of risk across divisions is approximately *100 times* greater than that of the price. One explanation as to why banks do not price discriminate geographically is *place-based discrimination* regulations, which prohibit practices that deny individuals in certain areas equitable access to credit through discriminatory pricing or underwriting. Such practices have redistributive effects: subsidizing lower-income borrowers and divisions that persistently underperform expectations.

References

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

Banerjee, A. V. and B. Moll (2010, January). Why does misallocation persist? *American Economic Journal: Macroeconomics 2*(1), 189–206.

Baqaee, D. R. and E. Farhi (2020, 02). Productivity and Misallocation in General Equilibrium*. *The Quarterly Journal of Economics 135*(1), 105–163.

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

Beck, T., R. Levine, and A. Levkov (2010). Big bad banks? the winners and losers from bank deregulation in the united states. *The journal of finance 65*(5), 1637–1667.

Becker, B. (2007). Geographical segmentation of us capital markets. *Journal of Financial economics 85*(1), 151–178.

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.

Cetorelli, N. and L. S. Goldberg (2012). Banking globalization and monetary transmission. *The Journal of Finance 67*(5), 1811–1843.

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. and V. Venkateswaran (2019, July). The sources of capital misallocation. *American Economic Review 109*(7), 2531–67.

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.

Loutskina, E. and P. E. Strahan (2015). Financial integration, housing, and economic volatility. *Journal of Financial Economics 115*(1), 25–41.

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.

Morgan, D. P., B. Rime, and P. E. Strahan (2004). Bank integration and state business cycles. *The Quarterly Journal of Economics 119*(4), 1555–1584.

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.

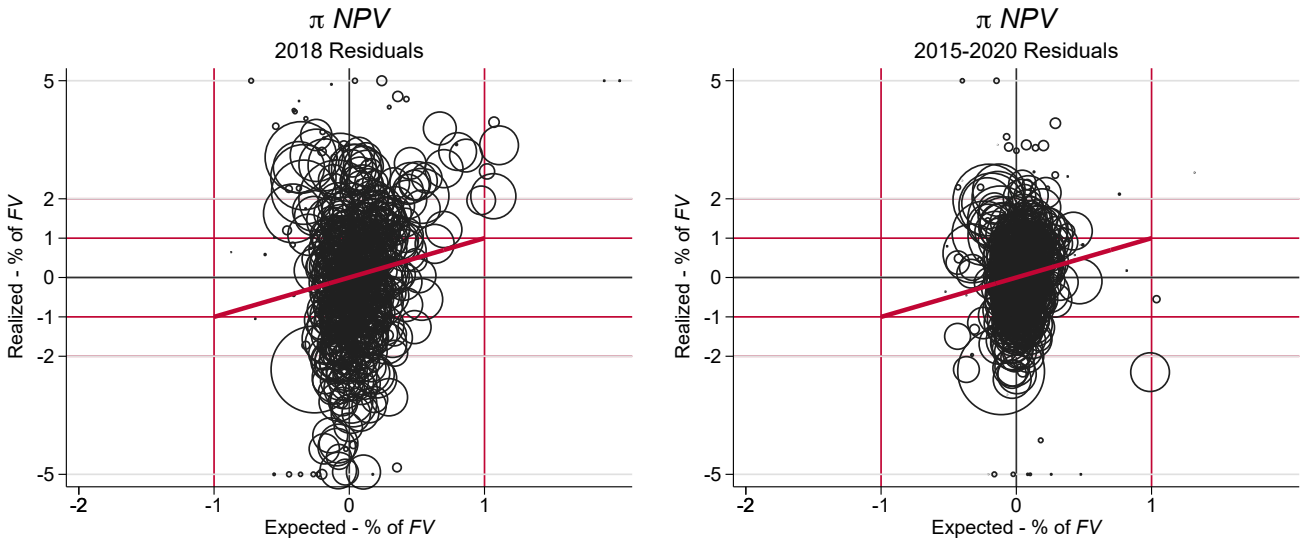
Whited, T. M. (2001). Is it inefficient investment that causes the diversification discount? *The Journal of Finance*.

Appendices

A Returns to Capital: Additional Results

The heterogeneity in divisional performance could be explained by the characteristics of loans (e.g., interest rate, size, maturity) originated across divisions. To isolate these differences, we extract residuals from a regression of loan performance on observable characteristics. Figure 2 displays the heterogeneity in aggregated division-level residuals of expected and realized performance after conditioning on loan and applicant characteristics. The figure shows a large dispersion for both expected and realized performance across all three measures. In unreported results, we also find a dispersion in the average product of capital exists evenwithin loans originated at the same price.

Figure 2: Returns Heterogeneity: Residuals



Note. Scatterplots for the division level expected and realized returns to capital. To residualize, we control for the credit score (9 bins, one for those will null score), the face value, term, and interest rate (4 bins each), and calendar year. We also include the credit score as a regressors to capture within-bin heterogeneity. We obtain residuals at the loan level (1 million plus). We then collapse these residuals to obtain division-level measures.

B Wedge and Loss Calculations: Additional Results

Table 9: Summary Statistics—Wedge Ingredients

Variable	Definition	Unit	N	Mean	SD	IQR	%
c	Cost of Funds	% APR	21	15.7	0.023	0.024	0.2%
R^s	Spread	% APR	21	3.2	0.049	0.021	0.7%
R	Rate	% APR	21	18.9	0.043	0.034	0.2%
T	Term	Months	21	29.0	0.28	0.30	1.0%
	Score		21	1409	15.8	9.0	0.6%
Θ	Realization	%	21	35.1	4.5	2.1	6.0%
Θ'	Slope		21	-12.2	3.0	3.1	25.4%
σ	Elasticity		21	6.8	1.1	1.9	27.9%
$\frac{\sigma-1}{\sigma}$	R coefficient		21	0.85	0.03	0.04	4.7%
q	Realization	% of FV	21	1.86	0.53	0.58	31%
q'	Slope		21	0.50	0.44	0.69	138%
δ	Elasticity		21	5.3	4.6	6.2	117%
$\frac{\sigma-\delta}{\sigma}$	q coefficient		21	0.21	0.71	1.04	495%

Note. L and H stand for rate cuts and hikes, respectively.

C Laws on Geographic Discrimination

C.1 United States (U.S.)

The Fair Housing Act (FHA) (42 U.S.C. §§3601-3619) prohibits redlining and lending discrimination based on location. The term "redlining" refers to the practice of mortgage lenders drawing red lines around specific areas or neighborhoods on maps to indicate their reluctance to make loans in those locations.

The Equal Credit Opportunity Act (ECOA) (Title VII of the Consumer Credit Protection Act of 1974) makes it *unlawful for any creditor to discriminate against any applicant* in a credit transaction *on the basis of race, color, religion, national origin, sex or marital status, or age*. While *place of residence* is not explicitly listed, U.S. regulators interpret ECOA, alongside the FHA, to prohibit redlining—denying or discouraging loans in certain neighborhoods based on their demographics. Regulatory enforcement ensures that lenders cannot consider the characteristics of a *neighborhood or geographic area* of a property (such as the race or ethnicity of its residents) in credit decisions. Regulations and enforcement guidance underscore that lenders may not consider the characteristics of the neighborhood or geographic area of a property (such as the race or ethnicity of its residents) in credit decisions. ECOA is enforced by agencies like the CFPB and DOJ; the DOJ may file lawsuits for patterns of lending discrimination, and lenders found engaging in geographic redlining face penalties including civil fines and corrective orders.

C.2 European Union (EU)

The Payment Accounts Directive establishes a right to basic banking services and explicitly bans geographic discrimination in that context. Article 15, titled "Non-discrimination," requires Member States to ensure that *credit institutions do not discriminate against consumers legally resident in the Union by reason of their nationality or place of residence* when those consumers apply for or access a payment account. In other words, an EU consumer cannot be refused a bank account (or given worse terms) solely due to living in a different EU country or region. Article 16 further guarantees the right to open a basic payment account regardless of the consumer's place of residence within the EU. National laws transpose these rules and designate authorities to supervise compliance. For example, a bank in one EU state offering accounts must accept applications from residents of other EU states under the same conditions. Enforcement can include regulatory sanctions imposed by national competent authorities for banks that unjustifiably reject customers based on address or nationality.

In addition, the Mortgage Credit Directive (MCD) enforces fairness in mortgage lending across EU member states; and the Geo-blocking Regulation prohibits unjusti-

fied location-based financial restrictions.

C.3 France

French law broadly forbids discrimination in the provision of goods or services, and it explicitly includes geographic criteria.

The French Penal Code (Article 225-1) lists "le lieu de résidence" (place of residence) among protected grounds alongside origin, sex, disability, etc. and hence criminalizes place-of-residence-based discrimination in lending. Consequently, a lender or bank in France that refuses credit, insurance, or a bank account to someone because of where that person lives is committing a criminal offense (assuming no legitimate risk justification). The law's enforcement has teeth: under Article 225-2, such discrimination is punishable by imprisonment and monetary fine for individuals, and higher fines or penalties if, for example, the discrimination occurs in a place open to the public or by a person in a position of authority. These provisions were used, for instance, to combat *banlieue discrimination* where residents of certain poor suburbs were unfairly denied financial services. French regulators and the Défenseur des droits (ombudsman) can investigate complaints; egregious cases are referred for prosecution, while administrative measures (like warnings or injunctions) can address lesser infractions. This robust framework ensures that geographic location alone cannot be used as a basis to refuse banking services in France.

In addition, the Housing Code ensures equal access to housing-related credit; and the French Competition Authority prohibits financial market segmentation by region.

C.4 Turkey

Law on Human Rights and Equality The Law No. 6701 prohibits both public and private entities from discriminating in offering services, which includes banking and finance. The law's Article 3(2) enumerates protected grounds: sex, race, colour, language, religion, belief, sect, philosophical or political opinion, ethnic origin, wealth, birth, marital status, health status, disability and age. While place of residence is not explicitly named here, many forms of indirect geographic discrimination are covered. For example, refusing loans in a predominantly Kurdish region could constitute ethnic origin discrimination; denying services in poorer districts might be seen as discrimination based on wealth or socio-economic status (which is a listed ground in Turkish law). The Equality law created an independent enforcement agency (TİHEK) empowered to investigate complaints and apply sanctions. If a bank were found to have engaged in unjustified geographic bias against consumers, TİHEK could impose administrative fines.

Additionally, Article 5 of Law 6701 mandates equal treatment in access to "services... provided by private law persons," making it clear that banks must serve customers without unjustified differentiation by region. In sum, although Turkish law stops short of naming "residence" as a standalone protected category, it polices geographic discrimination indirectly through other grounds, and regulators can take action if, say, a pattern suggests a certain area's residents are being excluded due to who they are or their economic condition. Compliance is further reinforced by Turkey's Constitutional equality clause (Article 10) and its commitments under international human rights treaties that demand nondiscrimination in economic life.

In addition, the Consumer Protection Law (Law No. 6502) ensures equal treatment of consumers; the Law on the Protection of Competition (Law No. 4054) restricts anticompetitive geographic pricing practices; the Banking Law No. 5411 enforces transparency in lending; and the Turkish Commercial Code (Law No. 6102) prevents unfair geographic price differentiation.

Table 10: Laws on Geographic Discrimination

Country	Law Name	Citation
U.S.	Equal Credit Opportunity Act (ECOA)	*15 U.S.C. §§ 1691–1691f*
	Fair Housing Act (FHA)	*42 U.S.C. §§ 3601-3619*
	Community Reinvestment Act (CRA)	*12 U.S.C. §§ 2901-2908*
	Dodd-Frank Act (UDAAP Provisions)	*12 U.S.C. §§ 5301–5641*
	Sherman Act (Antitrust)	*15 U.S.C. §§ 1-7*
	Robinson-Patman Act	*15 U.S.C. §§ 13-13b, 21a*
E.U.	Race Equality Directive	*Council Directive 2000/43/EC*
	Unfair Commercial Practices Directive	*Directive 2005/29/EC*
	Mortgage Credit Directive (MCD)	*Directive 2014/17/EU*
	Geo-blocking Regulation	*Regulation (EU) 2018/302*
	Payment Accounts Directive	*Directive 2014/92/EU*
France	Penal Code (Discrimination Law)	*Code pénal, Article 225-1*
	Housing Code	*Code de la construction et de l'habitation*
	French Competition Authority Guidelines	*Various*
Turkey	Law on Human Rights and Equality	*Law No. 6701*
	Consumer Protection Law	*Law No. 6502*
	Law on the Protection of Competition	*Law No. 4054*
	Banking Law	*Law No. 5411*
	Turkish Commercial Code	*Law No. 6102*