Financing Small Business: Competition Effects and Welfare Gains from Credit Assistance*

Maggie Isaacson[1] Chinmay Lohani[1]

Nov 2025

Updated Regularly. **Link to Latest Version**

Abstract

Governments operate large-scale credit assistance programs to promote the entry and survival of small firms. We study their effects on welfare and competition by analyzing the Small Business Administration's 7(a) program—the largest credit guarantee scheme in the United States. We develop and estimate a model of firm behavior with credit constraints using confidential Census microdata from the hotel industry, the program's largest recipient. Counterfactual analysis reveals that the program expands credit supply to inefficiently credit-constrained firms. We estimate that 7(a) guarantees cost $24 million and raise total welfare in targeted hotel markets by $54 million. Of these gains, $31 million come from valuable firms that fail to operate without guarantees, and $23 million from enhanced competition that lowers prices and expands markets. To understand how targeting can improve program design, we examine alternative allocations of 7(a) awards. Directing funds to concentrated markets increases consumer welfare by promoting entry, while awards in larger, less concentrated markets primarily benefit hotels by reducing exit and borrowing costs.

*Maggie Isaacson: University of Pennsylvania, Chinmay Lohani: University of Pennsylvania. We are indebted to Juan Camilo Castillo, Ulrich Doraszelski and Aviv Nevo for their advice on this project. We thank Francesco Agostinelli, Juan Pablo Atal, Winston Dou, Daniel Garrett, Michael Lee, Guillermo Ordonez, Asani Sarkar, Luke Taylor, Daniel Waldinger and the seminar participants at JHU, Penn, NYU Applied Microeconomics PhD Student Conference, the Federal Reserve Bank of New York and at the Penn Applied Micro Lunch for their helpful comments. We gratefully acknowledge support from the School of Arts and Sciences at the University of Pennsylvania and the New York Federal Reserve. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3109. (CBDRB-FY25-P3109-R12684). Estimates cleared for disclosure by U.S. Census under disclosure ID 12684, 12761. All remaining errors are our own.

1 Introduction

Credit frictions distort firm entry and exit, with direct consequences for competition and welfare.[1] These distortions are especially acute for small businesses, which rely on debt and systematically face binding credit constraints.[2] In key service industries, these small firms account for a majority of revenue.[3] When credit frictions suppress their entry or accelerate their exit, market structure shifts—and with it, the distribution of gains.

Recognizing credit frictions as a policy concern, governments have intervened at scale through credit assistance schemes. Between 2010 and 2019, the world's ten largest economies supported over $1 trillion of small business loans through credit assistance programs.[4] These programs often cite competition as a central objective: the U.S. Small Business Administration (SBA) seeks to "preserve free competitive enterprise," while India's NITI Aayog[5] describes small firms as drivers of "vigorous competition." Yet whether such programs deliver on that promise remains unclear. If assisted credit reaches financially constrained and surplus creating firms, it may enhance competition and lower prices. But if it simply subsidizes firms that steal business from competitors without generating additional surplus, such programs may induce excessive entry with little social benefit—concerns that are reflected in ongoing Congressional scrutiny of SBA lending.[6]

In this paper, we study the welfare effects of credit assistance to small businesses. We focus on credit allocation from loan guarantees, and examine their effects on market expansion and price competition. Our empirical setting is the SBA's 7(a) program—the largest credit-assistance scheme for U.S. firms—between 2010 and 2019. The 7(a) program provides loan guarantees to lenders and primarily targets service firms. We begin by showing that the program increases firm entry and reduces exit, using loan award data linked to firm-level Census micro-data. To evaluate the program's welfare consequences,

[1] Midrigan and Xu (2014)

[2] In the 2020 Small Business Credit Survey (fielded pre-pandemic), only 20% of financially healthy small firms reported having sufficient cash reserves to maintain normal operations during a two-month revenue loss. Among the remainder, 17% indicated they would have to close if unable to manage the shock. See Federal Reserve Banks (2020), *Small Business Credit Survey: 2020 Report on Employer Firms* Separately, the OECD (2019) reports that new and young firms face systematically more restricted access to finance than established firms, particularly for fixed investment and working capital. See OECD (2019), *SME and Entrepreneurship Outlook*, Chapter 6. (Mills, McCarthy, and Mayer 2014)

[3] Small firms' share of revenue in the U.S.: 56% in accommodation and food services, 54% in professional services, 82% in personal care.

[4] Authors' calculations, see Appendix 9.1

[5] The Government of India's premier policy think tank and planning body.

[6] Dilger (2019) and U.S. Small Business Administration Office of Inspector General (2019)

we develop a model of service-sector firms with entry, exit, and financial frictions. In our model, we explicitly micro-found lender behavior and depart from reduced-form descriptions of loan supply (Adams, Einav, and Levin 2009; Crawford, Pavanini, and Schivardi 2018; Hennessy and Whited 2007). This structure allows us to conduct counterfactual analyses of credit policy, quantify the value of current 7(a) awards, and assess alternative program designs.

We begin our analysis by examining the 7(a) program's direct effects on firm entry and exit. To study exit, we implement a matched difference-in-differences design. Awarded firms are matched to controls within industry-age-award-year bins that have a similar propensity to receive 7(a) assistance, based on revenue, employment, and growth. This allows us to compare treated firms to control firms with similar *ex ante* probability of receiving 7(a) awards, but different award status. We find that 7(a) awards reduce the probability of exit by nearly one percentage point over a five-year horizon—a 22% decline relative to the baseline exit rate.

Having analyzed effects on exit, we next examine how the program affects firm entry. We use cross-market variation in 7(a) loan volume to study how increases in award amounts affect new business formation. To address endogeneity in lending, we follow Greenstone, Mas, and Nguyen (2020) and construct a shift–share instrument. The instrument interacts each market's pre-period bank exposure with national changes in bank lending across years. This approach isolates quasi-exogenous changes in city-level 7(a) lending driven by historical bank presence. Banks experience idiosyncratic shocks to lending, which creates variation across cities based on bank footprint. We estimate that a 10% increase in 7(a) loan volume leads to a 1.7% increase in firm entry.

To assess the welfare implications of increased entry and reduced exit from 7(a) assistance, we develop a model. The model features two markets: an accommodation market, where guests purchase hotel services, and a debt market, where hotels demand loans. These markets interact: hotels compete in the accommodation market and earn profits, which they use to secure debt and finance their operations. In turn, debt allocation determines entry and exit of hotels, which shapes the competitiveness of accommodation markets.

In the accommodation market, competition determines hotel profitability. Hotels sell rooms to guests and generate profits. Their earnings vary across years depending on demand conditions, operating costs, and the degree of local competition. These profits serve as the basis for loans, as hotels pledge future income to obtain financing.

On the demand side of the debt market, hotels borrow to finance entry and ongoing operations. When a hotel enters, it takes on debt to cover its entry cost. Part of this cost is real estate expenses, financed through a fully recoverable loan that is repaid via stable mortgage payments. The remaining expenses—furniture, equipment, advertising, and cash reserves—are financed through a loan that carries an interest rate reflecting its risk. Once a hotel begins operating, it repays its debt from its profits. If profit falls short, the hotel demands additional loans to roll over its liabilities. This is only possible until it reaches a maximum level of debt, namely the debt limit, beyond which lenders refuse further credit. When a hotel is unable to obtain necessary funds, it defaults and exits the market.

Lenders form the supply side of the debt market and determine the allocation of credit. They are perfectly competitive and set interest rates based on default risk. This risk is costly because lenders only recover a fraction of their loan when borrowers default, called the recovery rate. Anticipating these losses, lenders charge higher interest rates to compensate for greater risk. Consequently, larger loans, which are less likely to be repaid, carry higher rates. Beyond a point, higher rates can no longer offset default risk, and borrowers become ineligible for additional credit. This point defines their debt limit.

The model's key friction arises from three features: hotels rely entirely on debt for financing, lenders recover only part of their loans in default, and profits are volatile. Because hotels borrow only through debt, lenders cannot vary repayments across different profit outcomes. When profits are low, some firms default and lenders incur losses. Lenders raise interest rates to cover this risk, but higher rates also increase the likelihood of default, preventing full compensation. As a result, credit supply is limited.

Loan guarantees expand credit supply by altering lenders' recovery rate. The SBA 7(a) program lowers lenders' losses by providing partial guarantees in the event of default. A higher recovery rate makes loans less risky for lenders. This expands debt limits and lowers interest rates for hotels, making them more profitable. This encourages entry in markets, and enables firms to survive longer. As a result, accommodation markets change, shifting prices and welfare.

We estimate our model in stages, beginning with consumer demand. For this purpose, we combine Census microdata on firm revenue with a collected dataset of hotel prices. Using a nested logit demand system, we capture consumer price sensitivity and substitution patterns among close alternatives within four nests—hotels, motels, other accommodations and an outside option. We address price endogeneity using competitor characteristics as instruments. Our estimates imply that consumers are highly price sensitive: the median own-price elasticity across hotels is –5.1, with a mean of –5.4 . The estimated nest parameter is 0.65, which implies significant within-group substitution and business-stealing effects.

Next, we estimate the parameters governing credit markets and firm dynamics. The key parameters include the degree of profit uncertainty, which determines underlying volatility, the magnitude of entry costs, which shape entry barriers and firm indebtedness, and the recovery rate, which determines lenders' losses in default. We obtain the average recovery rate from 7(a) loan-level default data, which show that lenders recover 61% of principal upon default. We identify entry barriers by observing patterns of entry in the data. The model then links profit volatility to firm exit through its effect on credit limits. This allows us to identify features of the profit stream from exit data alone. Thus, we estimate our model using economic data on entry, exit, and revenue. This feature is crucial for our setting where, as is typical, rich firm-level financial data are unavailable for small businesses (Brown et al. 2024).

Our estimates indicate substantial profit uncertainty: the standard deviation of demand-driven profit changes equals roughly 37% of mean firm-level profits. Operating cost shocks are large and account for higher exit among smaller firms—-for firms in the bottom 5% (by profit), the standard deviation of operating costs is nearly 60% of mean profits. Low recovery rate and volatility imply that firms can borrow between three to four times their annual variable profit as working capital. Accounting for variable costs, the estimates imply that firms earning less than $38,000 in annual operating profits fail to break even in expectation.

Using the estimated model, we evaluate the program's welfare effects. By comparing outcomes in counterfactuals with and without 7(a) assistance, we find that credit frictions significantly restrict efficient firm entry and survival across our sample hotel markets. The 7(a) program expands credit supply and increases total welfare by $54 million. These gains come at a fiscal cost of $24 million in guarantee payouts, implying a positive net benefit and a return of $2.25 per dollar of public expenditure.

We decompose the total welfare gains into two components: gains from a larger set of options and from stronger competition. Of the total welfare increase, $31 million reflects value added by the presence of additional options in the accommodation market. On net, 7(a) awards increase credit allocation towards valuable hotels that expand the scale of the market, but were previously credit constrained. About $23 million reflect gains from increased competition. More firms enter and remain active, competition intensifies, and prices fall—by 0.6% on average, and up to 23% in the most concentrated markets. Cheaper prices create substantial benefits for price sensitive consumers in these markets.

We now turn to the incidence of 7(a) program benefits. Consumers are the largest beneficiaries, gaining \$46 million—nearly half from lower prices, and the remainder from having a larger set of hotels to choose from. These gains are geographically concentrated: the most concentrated decile of markets, which represents 8% of revenue, account for over 20% of total consumer surplus. Hotel surplus increases by approximately \$8 million, but with significant heterogeneity across markets and firms. Marginal entrants—hotels that enter only because of the program—gain surplus at the expense of incumbents due to significant business stealing effects. After accounting for hotel entry costs and exit by incumbents, net producer surplus declines slightly in markets with new entry. Hotels that receive 7(a) awards, but are not marginal entrants, create welfare cost-effectively. A 7(a) award lowers their borrowing costs and extends firm survival. This credit expansion does not incur large fiscal costs, since awards reduce exit probability and associated guarantee payments, partially offsetting the cost of credit assistance.

We then examine how program targeting can improve welfare outcomes. As a simple exercise, we shift 7(a) awards to more concentrated markets, measured by HHI.[7] Specifically, we remove awards from the least concentrated third of markets, and double awards in the most concentrated third. Relative to a no-7(a) baseline, this new allocation increases consumer surplus by \$52 million and producer surplus by \$5 million, at a fiscal cost of \$30 million, resulting in a lower multiplier (relative to current 7(a) awards) of nearly two. These results are consistent with our earlier finding that consumer gains are concentrated in less competitive markets, while producer gains from new entry remain limited due to substantial business-stealing effects.

Finally, we consider an optimal allocation of program funds under different welfare weights. In our simple exercise, we divide market-firm combinations into a two dimensional grid by profit and HHI. The government allocates awards within profit by HHI bins with zero or one probability. We consider two scenarios—one where government values consumer and producer surplus equally, and one where the government only values consumer surplus. Under either scenario, the optimal award scheme avoids the smallest 20% of hotels. While they can increase surplus, especially in concentrated markets, these hotels have a higher exit rate since they are substantially affected by operational costs, especially in their early, highly indebted years. Any gains created over their short life-span are not cost effective, and excluding them significantly increases the fiscal multiplier.

We find that when consumer and producer surplus are equally important, the optimal

[7] Herfindahl–Hirschman Index

allocation awards mid-profit firms—in the 20th–80th percentiles—in concentrated markets with many hotels. In these markets, 7(a) awards rarely induce marginal entry, almost exclusively in the wake of firm exit (new entry in markets with many firms is atypical, similar to Bresnahan and Reiss (1991)). Instead, 7(a) awards decrease hotel borrowing costs, increase hotel dividends and decrease exit. In the absence of new entry costs, and reduced fiscal costs due to exit, the program is cost-effective. Under this award scheme, consumer welfare rises by $45 million, hotel welfare by $130 million, and the government pays $40 million in guarantees—implying a return per fiscal dollar above four.

Under a consumer-focused weighting, the optimal policy shifts towards hotels in the most concentrated markets, where 7(a) encourages new entry, lowers prices and expands consumer surplus. Producer gains, however, remain limited: business stealing erodes incumbent profits, and entrants' fixed costs offset part of their profits. Even so, removing support for the smallest 20% of hotels reduces guarantee payments and enhances cost-effectiveness. In this case, the policy yields $150 million in consumer surplus and $35 million in producer surplus at a fiscal cost of $45 million.

Taken together, our results highlight three broad insights for credit policy. First, credit constraints in service industries—such as the small hotel markets we study—can meaningfully distort market structure, with consequences for welfare. Second, credit guarantees like 7(a) awards can generate welfare gains cost-effectively by expanding the supply of previously constrained firms. Credit guarantees accomplish this by insuring lenders against default, which expands firms' access to capital, increases firm supply, and strengthens competition. However, these gains are uneven: non-beneficiary firms are hurt by increased competition and business-stealing effects. Finally, the effects of credit assistance depend critically on program targeting. Optimal program design therefore requires careful targeting aligned with the planner's welfare objective.

Related Literature We contribute to a large literature that studies how small business credit assistance programs affect firm outcomes. Loan-guarantee schemes have yielded mixed results: some papers find sales, employment and productivity gains (Banerjee and Duflo 2014; Bazzi et al. 2023; Brown and Earle 2017), while others document limited real impacts or higher default risk (Greenstone, Mas, and Nguyen 2020; Lelarge, Sraer, and Thesmar 2010). Related work on banks and guarantees studies how lenders mediate changes in debt markets (Bachas, Kim, and Yannelis 2021; Stillerman 2024). We extend this literature by quantifying equilibrium effects of loan guarantees on *downstream* market structure. Studying how the 7(a) program affects competition and expands hotel markets, we identify a substantial and previously unexplored benefit of loan guarantees.

More broadly, our paper connects to research showing that financial frictions shape economic outcomes. Foundational theories of debt, limited pledgibility and collateral point out the importance of financial frictions on markets (Hart and Moore 1994; Holmström and Tirole 1997; Kiyotaki and Moore 1997; Townsend 1979). Subsequent empirical work in macroeconomics demonstrates that credit influences investment, entry, and productivity (Buera, Kaboski, and Shin 2011; Cavalcanti et al. 2021; Midrigan and Xu 2014). Recent research emphasizes two key insights: much lending relies on cash-flow-based finance rather than fully secured collateral (Kermani and Ma 2020), and that recovery frictions—arising from asset specificity, distress-related degradation, and legal recovery costs—significantly limit credit supply (Benmelech and Bergman 2009; Dou, Wang, and Wang 2022; Eisfeldt and Rampini 2009; Ma and Scheinkman 2020). We contribute to this literature by showing that low recovery rates create a substantial financing wedge, and that that loan guarantees can narrow this wedge. Our structural model embeds recovery frictions directly, allowing us to quantify how SBA guarantees narrow this wedge and affect firm entry, survival, and market-level welfare.

We also relate to the literature at the intersection of finance and industrial organization. A rich literature of bank and credit-market competition examines how lenders set terms under screening and informational frictions (Adams, Einav, and Levin 2009; Crawford, Pavanini, and Schivardi 2018; Einav, Jenkins, and Levin 2012; Kastl 2012). On the other hand, the literature on the interaction of financial and product markets is relatively sparse. Early empirical work explored the role of financial distress and capital structure in explaining supermarket markups (Chevalier 1995; Chevalier and Scharfstein 1996). More recently, Hortaçsu et al. 2019 demonstrate how financial distress creates 'indirect costs' in used auto sales, while Fan, Kühn, and Lafontaine (2017) show how drops in collateral value leads to changes in business franchising by hotels. Despite this progress, empirical evidence on how financial access shapes market structure and competition remains limited. Doraszelski, Gomes, and Nockher (2022) is a notable exception that provides simulation-based evidence on industry dynamics under financial frictions. We contribute by providing the first empirical estimates of how credit policy affects market structure and competition. In particular, we show how financial markets shape firm entry and survival in service markets with direct consequences for consumer welfare.

Finally, our work contributes to the IO literature on market entry and related policies. Classic models show that business-stealing externalities can generate inefficient entry (Mankiw and Whinston 1986), and empirical studies show their welfare consequences, creating a role for policy to regulate entry (Berry and Waldfogel 1999). A natural question that arises then, is whether government policies are effective in directing entry towards welfare improving directions. While theoretical work examines involved trade-offs (Dixit and Kyle 1985; Pflüger and Südekum 2013), empirical evidence on entry subsidies and their effectiveness remains limited. Pesendorfer and Schmidt-Dengler (2003) and Dunne et al. (2013) evaluate such policies through counterfactual simulations, while Fan and Xiao (2015) provide empirical evidence from deregulated U.S. telephone markets. We extend this work by demonstrating that loan guarantees serve as an effective policy lever for influencing entry and exit in small-business dominated markets. Our analysis decomposes welfare effects into market expansion and business stealing, finding that market expansion dominates in small hotel markets. As a result, market expansion policies like loan guarantees create net surplus gains.

Outline The rest of the paper proceeds as follows. First, we describe the data and institutional setup in Section 2. Then, in Section 3, we give a stylized example explaining how credit frictions arise, and descriptive evidence that points to their presence in our setting. Afterwards, we show evidence of entry and exit effects of 7(a) awards in 4. Section 5 describes our empirical model, and Section 6 shows how the model is identified and estimated. We proceed with our counterfactuals in Section 7 and conclude in Section 8.

2 Data and Institutional Context

2.1 SBA 7(a) Program

The Small Business Administration (SBA) is a U.S. federal agency tasked with assisting small firm owners in establishing and maintaining their businesses. Our paper focuses on its flagship lending program, the 7(a) loan program, which is designed to expand credit supply to small businesses. It is a loan guarantee program, where private lenders provide loans and SBA guarantees a portion of the loan amount that is charged off (unpaid) by borrowers upon default. SBA 7(a) supported loans worth approximately $206 billion over more than half a million loan awards during our study sample, 2010-2019.

The program has some eligibility requirements: a 7(a) eligible business must be run for-profit, be independently owned and operated, and not be dominant in its field. It must also fall under the SBA's size requirements for small businesses, which vary by industry based on thresholds in either annual receipts or number of employees. Certain sectors, such as financial services, are excluded from the program. If a borrower is deemed eligible, the loan proceeds in several steps. The borrower approaches a lender with a financing request. The lender verifies eligibility and the borrower applies for a 7(a) loan. The SBA then reviews the application and makes a decision.[8]

The program serves as a lender of last resort through two key features. First, borrowers must exhaust all other financing sources, including personal equity, before qualifying for a 7(a) loan. Second, lenders must certify through the 'credit elsewhere' test that the loan would not be available on reasonable terms without the SBA guarantee. The SBA enforces this requirement through targeted lender reviews, with non-compliance potentially resulting in guarantee denials or enforcement actions from the Office of Credit Risk Management. Notably, while maintaining these strict requirements, the program explicitly prohibits lenders from denying credit solely due to insufficient collateral.

Once the loan is disbursed, the SBA intervenes only when the borrower defaults. First, the lender is required to pursue recovery—either by liquidating assets or enforcing personal guarantees. The SBA then covers the guaranteed share of the remaining charged-off amount. The SBA reports charge-offs to represent between $0.5\% - 2\%$ of award principal during this period. [9] (U.S. Small Business Administration 2024a)

2.2 Data

Our empirical analysis draws on three main sources: SBA 7(a) loan award records, confidential Census micro-data on firm outcomes, and a novel dataset of hotel prices collected for U.S. cities with a beneficiary hotel. In addition to these, we illustrate stylized facts about major industries that receive 7(a) support using IRS Statistics of Income (SOI) Corporation data. This section describes each data source and documents key patterns.

SBA 7(a) loan records identify beneficiary firms and provide rich information on loan terms and outcomes. These data are publicly available through Freedom of Information Act (FOIA) releases after 2016 and include the universe of loans disbursed since 1991. Each entry contains detailed information on borrower and lender identities, loan amounts, maturity, interest rates, SBA guarantee shares, and realized outcomes such as full repayment or default, along with defaulted amounts. These data allow us to track the flow of public credit to individual businesses and analyze variation in awards and outcomes across banks, borrowers, and time. We subset to awards between 2010 and

[8]Under SBA's Preferred Lenders Program, approved banks may approve and close 7(a) loans without prior SBA review. They must still confirm that the borrower and use of funds meet SBA rules and report the loan's main details—such as amount, purpose, maturity, and collateral—to the SBA so the agency can issue its guarantee and record the loan. In 2019, such loans represented 46% of loans by count, and 75.8% of loans by volume. (Dilger and Lowry 2020)

[9]Charge off rates decreased over time as per U.S. Small Business Administration 2024a.

2019 in our analysis to exclude large macroeconomic shocks outside of the scope of the project.

Our second data source is firm-level Census micro-data, which provides detailed information on firm inputs and outcomes. The Longitudinal Business Database (LBD) covers all U.S. employer businesses with EIN (Employment Identification Numbers) identifiers, providing annual data on employment, establishment counts, and firm entry and exit. We link these with revenue data from the Business Register, which provides us with revenues for a subset of these firms. Following Brown and Earle (2017), we link firm-level records to SBA loan data using employer identifiers, allowing us to compare assisted and unassisted firms along key margins.

To estimate consumer demand for hotels, which is the industry of focus for our structural analysis, we collect a dataset of hotel prices. For each city where at least one hotel received 7(a) assistance, we collect prices and observable hotel characteristics via web scraping. Specifically, we generate a latitude–longitude grid within each city and run targeted Google Maps queries. These searches return detailed information on each hotel's price, star category, average customer rating, and number of reviews along with a corpus of text reviews. We match establishments to known chains and brands to identify franchises. The appendix provides details on the scraping procedure, matching algorithm, and data cleaning steps.

Finally, we use IRS Statistics of Income (SOI) data to document cross-industry differences in firm characteristics and expenses. These data provide industry-level aggregates for all U.S. corporations filing income tax returns, including detailed categories of expenditures and debt holdings. This allows us to assess debt reliance at the industry level, extending beyond the firms and loans observed in our 7(a) sample. We use the SOI to compare debt dependence, collateralizable expenses, and profit volatility across industries targeted by the 7(a) program. These characteristics map directly into the debt restrictions emphasized in our model.

3 Stylized Example and Descriptive Patterns

We begin with a stylized example that illustrates how credit frictions arise in our setting. The example clarifies the role of three firm characteristics—debt reliance, incomplete recovery, and income volatility—in generating credit frictions.

Next, we show that industries targeted by the 7(a) program display these features in the data. We also show that, even within these industries, 7(a) loans support marginal borrowers. We then describe key features of the hotel industry that make it an ideal case study, including its exit patterns, volatility, and the small local markets that the program primarily reaches. Together, these patterns motivate the subsequent analysis.

3.1 Credit Frictions: A Stylized Example

We now describe a simple example that illustrates how financial frictions arise in our setting. There are three key ingredients that result in credit frictions: debt exclusive financing, low debt recovery and income volatility.

Consider a firm that operates for one period and requires a \$14 working capital loan to start. At the end of the period, the firm either either \$20 (good state) or \$10 (bad state) in profit with equal probability. Its expected payoff is \$15, so the firm is profitable net of its cost.

Assume that the firm can only seek debt financing. That is, it cannot sign contracts where payments depend on the realized profit. A lender provides loans and collects matured debt (principal + interest, if any) after the firm finishes operations. If the firm cannot pay its matured debt at the end of the period, it defaults. In the event of default, the lender recovers only \$6, losing the rest in recovery costs.[10]

The lender is perfectly competitive and must break even in expectation. To earn \$14 on average, the lender needs to charge an interest rate that gives it more than \$14 in the good state to compensate for earning strictly below \$14 in the bad state. The highest amount that the lender can recover in the good state is \$20, otherwise the borrower would default in all states.[11] By setting loan terms so that repayment in the good state equals \$20, the lender obtains its maximum expected return:[12]

$$\frac{1}{2}(20) + \frac{1}{2}(6) = \$13$$

which falls short of the \$14 required to finance the firm. Hence, the lender refuses to extend the loan—even though the firm is economically profitable.

Note the role of each of the three ingredients in generating credit friction. First, debt contracts fix payments across states, preventing repayment from varying with realized profits.[13] Second, weak collateral means that if the firm defaults, the lender recovers a

[10]This can be thought of as the scrap value of the business for the lender, though we abstract away from what determines losses in recovery.

[11]This setup parallels the one in Townsend (1979) of debt as an optimal contract, except we abstract from moral hazard: borrowers always repay in good states, so credit limits are higher.

[12]This is better than setting \$10, which always pays \$10

[13]This encodes moral hazard by borrowers. While their credit allocation would increase if they followed

fraction of its principal. Third, profit volatility creates dispersion around expected value, introducing states in which the firm cannot repay. Together, these elements make credit supply sensitive to recovery risk, constraining finance even for projects with positive expected value.

This scenario mirrors real-world service businesses, where firms must finance expenses such as furniture, equipment, and marketing through debt, and pledge to pay it through cash flow (Kermani and Ma 2020). Heavily collateralized loans are secure, and serviced easily. But if the owner possesses weak collateral, loans become infeasible, since imperfect recovery makes lenders reluctant to extend credit. These are precisely the type of constraints targeted by the 7(a) program.

Now suppose the government guarantees an additional \$2 upon default. The lender's \$20 dollar on maturity debt now collects an additional \$2 in the bad state, making its value:

$$\frac{1}{2}(20) + \frac{1}{2}(6+2) = \$14,$$

which matches the required startup loan. The government pays for the guarantee only in the bad state, so its expected cost is \$1. The guarantee enables a project that generates \$1 in private surplus at a fiscal cost of \$1.

Whether this intervention improves welfare depends on the social value of entry. If the firm merely steals profits from competitors, the \$1 surplus reflects reallocation with no added social surplus. On the contrary, the firm's social value may exceeds its private profit and further justify intervention. In the above example, if the additional firm created price drops worth \$2 for consumers, the overall welfare gain from the firm is \$3— a fiscal multiplier of three.

We now turn to the data and present the main descriptive patterns that motivate our analysis.

3.2 Descriptive Patterns

3.2.1 Program Allocation and Recipient Characteristics

The distribution of 7(a) awards reveals a strong concentration in service industries. Among the 110 recipient industries[14], the five largest recipients—hotels and motels,

state contingent payment schemes, there is no ex-post commitment device to ensure this. It can also be thought of as a technological constraint: monitoring technology is expensive relative to the scale of these loans.

[14] We use 3 digit NAICS code as the industry classification for this exercise

restaurants, gas stations with convenience stores, dental offices, and child care services—account for nearly a quarter of all loan volume. Among the top fifteen recipient industries (Appendix Table 9.2.1), fourteen are in the service sector.

Firms in these service industries share core features associated with credit constraints. They operate with higher leverage, spend a larger share of revenue on non-collateralizable inputs (e.g. labor, rent, and advertising), and experience more volatile cash flows than average. (see table (1)). These features suggest elevated risk (Kermani and Ma 2020) making such firms less attractive to lenders. The concentration of 7(a) guarantees in these sectors is thus consistent with the program's objective of addressing frictions arising due to debt reliance and imperfect recovery.

Table 1: Financial Characteristics of Top-5 SBA Industries vs. U.S. Economy

	Top-5 SBA Industries (Weighted Avg.)	U.S. Economy (Weighted Avg.)
Debt-to-Assets	40.8%	28.1%
Volatility	3.55	0.19
Non-Collateralizable Expenditures (% of Deductions)	30.5%	23.8%

Note: Data from IRS Statistics of Income Corporation Source Book (2019). Volatility is the weighted average of income coefficient of variation (CoV) for businesses with positive net income across industries. Non-collateralizable expenditures include compensation, rent, advertising, repairs, and other operating expenses. Top-5 SBA industries comprise Hotels, Restaurants, Gas Stations, Dental Offices, and Child Care Services, weighted by industry revenue.

3.2.2 Loans Are Marginal

While firms in targeted sectors are likely to be financially constrained, were the specific loans delivered by 7(a) marginal? We present suggestive evidence using 7(a) loan award data through a simple counterfactual exercise: we remove guarantees from a subset of loans and compare lender margins with and without these guarantees. Profitable loans should, on average, yield positive margins—otherwise lenders would be unlikely to provide them. If 7(a)-supported loans are truly marginal, average margins should be too low to justify lending in the absence of guarantees. We test this prediction using outcomes for matured loans.

Figure 1: Bank Margins with and without 7(a)

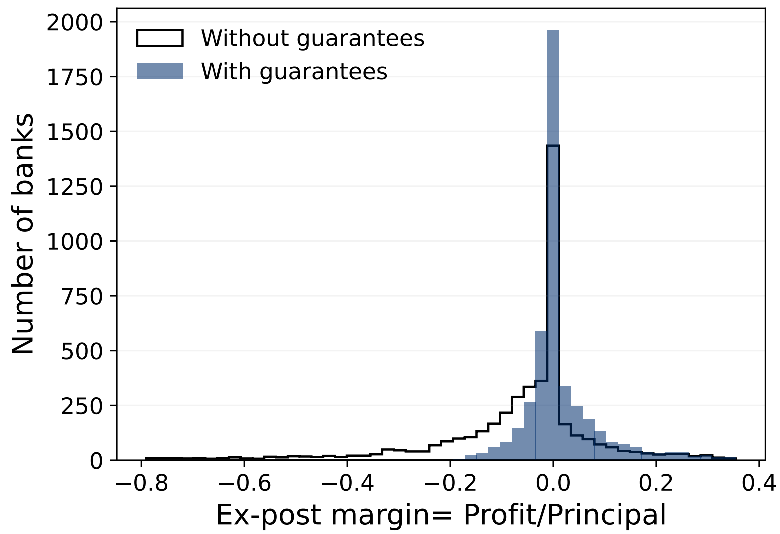

Note: This figure plots a histogram of bank margins from 2000–2019 on 7(a) loans. Here the margin is the inflation-adjusted payments from interest on principal, minus default losses, divided by the sum total of loan principal (2019 dollars). Shaded histogram represents factual margins, while unshaded bins represent counterfactual margins without guarantees.

We focus on loans that matured before 2019 to observe complete repayment histories. For these loans, we calculate each lender's accounting margin— inflation-adjusted interest income minus charge-off losses, per dollar of principal originated—under two scenarios: with and without guarantees. This sample skews toward shorter-maturity loans, which tend to be less collateralized and riskier, but avoids bias from censored outcomes. Figure (1) plots a histogram from our estimates.

Without guarantees, the average margin of all loans (total profits- total losses/total principal) is –6.9%. With guarantees, the distribution shifts rightward, eliminating severe losses and raising the average margin to 1.2%. This shift—from unprofitable to mildly profitable—is consistent with the program's core function of serving marginal loans.[15]

[15]Note that, this is counterfactual does not account for changes in borrower behavior in absence of guarantees. If these small loans are used for riskier projects *because* they are guaranteed, we would expect returns to be low because of excessively risky actions by borrowers. However, smaller 7(a) loans tend to be used for working capital expenses and to adjust for seasonal cash flows, rather than investments.(SBA 7(a) Loans 2025)

3.2.3 Empirical Focus: Hotels

We focus on hotels for our structural analysis for three reasons—substantial 7(a) program support, localized and well-defined markets, and distinctive financial characteristics. First, hotels are the largest recipients of 7(a) guarantees, accounting for $1.72 billion—or 8.4% of program volume—from 2010 to 2019, making them a significant part of the program's portfolio. Second, hotels operate in geographically distinct markets, providing variation across many independent markets and a tractable setting for measuring welfare through competition. Third, their financial characteristics—high leverage, earnings volatility, and reliance on non-collateralizable inputs—make them especially susceptible to credit frictions (Table 2). These features make hotels a natural case study to understand how credit frictions, and guarantees to alleviate such frictions, can change market structure and welfare.[16]

Table 2: Financial Metrics: Hotels vs. Broader Service Economy

	Hotels	Accom. & Food Services	Service Economy (Weighted Avg.)
Volatility	2.2	2.4	1.18
Debt Reliance (Debt-to-Assets)	50.2%	36.5%	22.1%
Non-Collateralizable Expenditures	63.7%	52.9%	49.1%

Note: Volatility is measured as the coefficient of variation of revenue receipts in IRS SOI (2019) weighted by industry revenue. Debt reliance is average debt-to-assets from IRS SOI (2019). Non-collateralizable expenditures include depreciation, repairs, advertising, and "other deductions" (IRS SOI (2019)).

Looking directly at 7(a) loans to hotels, we find patterns consistent with working-capital and residual property-financing needs. The average loan size is $1.7 million—well within the typical gap hotels face after exhausting conventional mortgage capacity and owner equity (see Appendix 9.5). These loans are heavily guaranteed, with an average SBA guarantee of 75% per loan. To study repayment outcomes, we examine hotel loans originated between 2001 and 2019, providing a large sample that includes longer-maturity loans. Among loans that have fully matured, 13.7% default, and lenders recover only about one-third of principal on defaulted loans. These high charge-off rates suggest that SBA guarantees play a critical role in enabling financing for these working-capital

[16] Although this focus limits external validity, detailed financial data at scale are scarce. We therefore study an industry with significant policy variation and where credit constraints are likely to bind, making their effects empirically detectable.

and secondary property expenditures that would otherwise be difficult to fund privately. (Office of the Comptroller of the Currency 2014a)

Typical Cities Affected by 7(a) To study the competitive context of 7(a)-backed hotel entry, we construct geographic accommodation markets using our scraped dataset of hotels and locations. Starting from 2,579 cities and towns where hotels were awarded 7(a) loans, we define 2,679 markets. Most towns (2,513) are treated as stand-alone markets. For 66 larger metropolitan areas (e.g., New York, Philadelphia), we subdivide them into 166 finer submarkets using k-means clustering on hotel coordinates, augmented with Google Maps neighborhood boundaries. This procedure isolates economically meaningful competitive boundaries in both rural and urban settings and serves as the basis for our demand estimation.

These markets are typically small. The median market contains seven hotels, and nearly half (1,217) contain fewer than five. Average nightly rates are $125, skewed up from a median of $99. One-third of these properties are classified as motels, and the median rating is two stars. Roughly 49% are franchised. These patterns confirm that 7(a) lending predominantly supports limited-service properties in small, less competitive markets. Whether entry in such markets improves welfare or simply redistributes surplus depends on substitution patterns and price responsiveness—questions we evaluate in our structural model. Additional summary statistics appear in Appendix Table (8).

3.2.4 Exit Patterns in Hotels[17]

We now examine patterns of firm exit among hotels. Using the Longitudinal Business Database (LBD), we document that exit rates[18] among hotels are substantial. Between 2010 and 2018, more than 20% of hotels in the sample exited, with rates exceeding 40% among the smallest firms, as shown in figure (2). These patterns are consistent with the financial fragility suggested by hotels' high leverage, revenue volatility, and reliance on non-collateralizable inputs.

To assess the role of volatility on firm exit directly, we focus on hotels entering between 2010 and 2014 and examine the relationship between within-firm revenue volatility and five-year exit. Volatility is measured as the coefficient of variation of annual revenues.

[17] Figures and estimates in this section are placeholders, and guided by public data (Hollenbeck 2017; U.S. Bureau of Labor Statistics, Business Employment Dynamics (BLS BDM) n.d.). Accurate estimates are withheld, pending Census disclosure review.

[18] This follows Davis and Haltiwanger (1992) in defining exit happening during the last year a firm appears in the LBD.

Figure 2: Exit Probability in Hotels

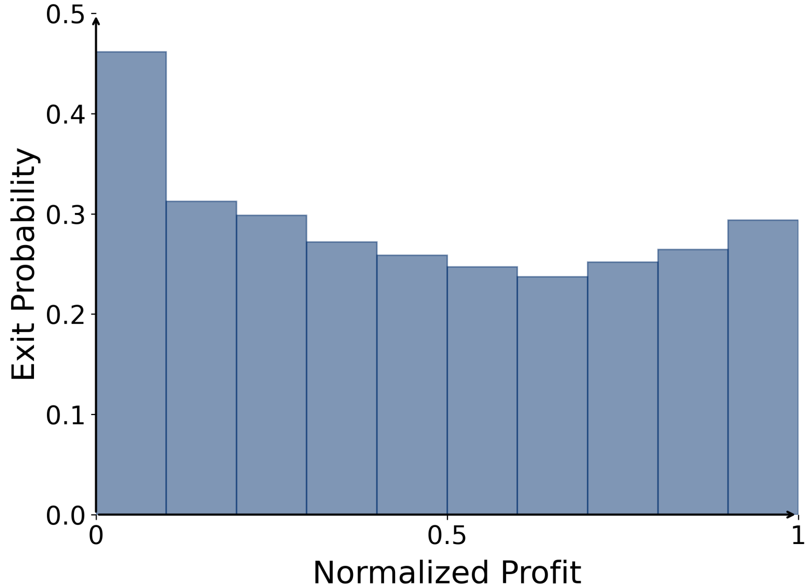

Note: This figure plots the probability of exit between 2010-2019 for hotels. Estimates cleared for disclosure by U.S. Census under disclosure ID 12761.

The average volatility in the same is 0.12. Table (3) shows a strong positive association: an increase in standard deviation of revenue by 10% of its mean predicts a 2.3 percentage point increase in exit probability. The relationship remains robust after controlling for firm size (via revenue bin fixed effects), age, and industry-year effects, suggesting that volatility is a meaningful predictor of exit risk.

In summary, the hotels operate in a volatile industry characterized by high turnover, where they need to finance expenses that are difficult to collateralize. SBA 7(a) guarantees appear to provide residual financing with substantial guarantees that support marginal firms, where conventional credit markets fail. These loans are concentrated in smaller hotel markets, where entry may reshape local market structure.

We now assess whether 7(a) guarantees measurably affect firm dynamics—entry and exit—using quasi-experimental variation in program exposure.

4 Evidence of the 7(a) Program's Effects on Entry and Exit

We examine whether 7(a) awards affect business survival and formation. For this, we employ two quasi-experimental approaches: a matched difference-in-differences design for exit effects and a shift-share instrumental variable strategy for entry effects.

Table 3: Revenue Volatility and Firm Exit

	(1)	(2)	(3)	(4)
Revenue Volatility (CV)	0.210	0.212	0.189	0.180
	(0.018)	(0.019)	(0.019)	(0.015)
Industry × Year FE		Yes	Yes	Yes
Revenue Bin FE			Yes	Yes
Age FE				Yes

Note: OLS estimates for hotels entering 2010-2014. Revenue volatility is the coefficient of variation of annual revenues. Exit is defined as the last year a firm appears in the LBD (2010-2018, excluding COVID years). Industry is 3-digit NAICS. Revenue bins are deciles of average firm revenue. Standard errors clustered at city level. Estimates cleared for disclosure by U.S. Census under disclosure ID 12761.

4.1 Exit Effects: Matched Difference-in-Differences

We estimate the effect of 7(a) receipt on firm exit using a staggered difference-in-differences specification that compares treated firms to matched controls.

Treated firms are recipients of 7(a) loans between 2010 and 2019. For each treated firm, we select a matched control from the same 2-digit NAICS industry that (i) never received a 7(a) loan during the sample, (ii) is in a market where no firms in that industry received 7(a) during the sample (to preserve SUTVA), and (iii) has a similar ex-ante propensity to receive a loan based on observables. Matching is performed within firm-age groups (0–1, 2–5, 6–10, 10+ years) using pre-treatment averages and growth rates of revenue and employment; we implement one-to-one nearest-neighbor matching with a 0.1 s.d. caliper. (Rosenbaum and Rubin 1983, Heckman, Ichimura, and Todd 1998) Because matching aligns treated firms with controls of the same age and calendar year, we assign a pseudo-treatment year T_i to each control firm based on its matched treated firm, and define event time $t - T_i$ for both treated and control firms.

We estimate the effect of receiving a 7(a) award on the conditional probability of exit—that is, conditional on surviving for $k-1$ years after receiving an award, we compare the difference in exit probability in year k between treated and control firms. Our estimating equation is:

$$Y_{it} = \alpha_i + \lambda_t + \sum_{k \neq -1} \beta_k \, \mathbf{1}(t - T_i = k) \, D_i + \mathbf{X}_{it}' \gamma + \varepsilon_{it}, \tag{1}$$

where Y_{it} is an indicator for exit in year t; once a firm exits, it is removed from the panel.

D_i indicates whether firm i ever received a 7(a) loan. The coefficients β_k measure the difference in exit probability between treated firms and matched controls at event time k, normalized to zero in the year before treatment ($k = -1$). We include firm fixed effects (α_i) and calendar-year fixed effects (λ_t), and control for lagged revenue and employment in $\mathbf{X}_{it}$. Because we condition on survival through treatment, we cannot directly test for parallel pre-trends in exit. However, we validate the design by verifying balance in pre-treatment levels and trends of revenue, employment, and wages between treated and control firms.[19]

Event–study estimates indicate that the effect of 7(a) loans on firm survival has immediate effects (see figure 3). In the treatment year, the conditional probability of exit falls by roughly 2.1 percentage points from a 4.4 percentage point baseline—about a 47% decline. The effect attenuates over subsequent years. Two years after receiving the award, treated firms are as likely to exit as control firms. Heterogeneity by age shows that the response is concentrated among young firms (2–5 years old), which face the highest baseline hazard rates.

Figure 3: Effect of 7(a) Awards on Exit Probability

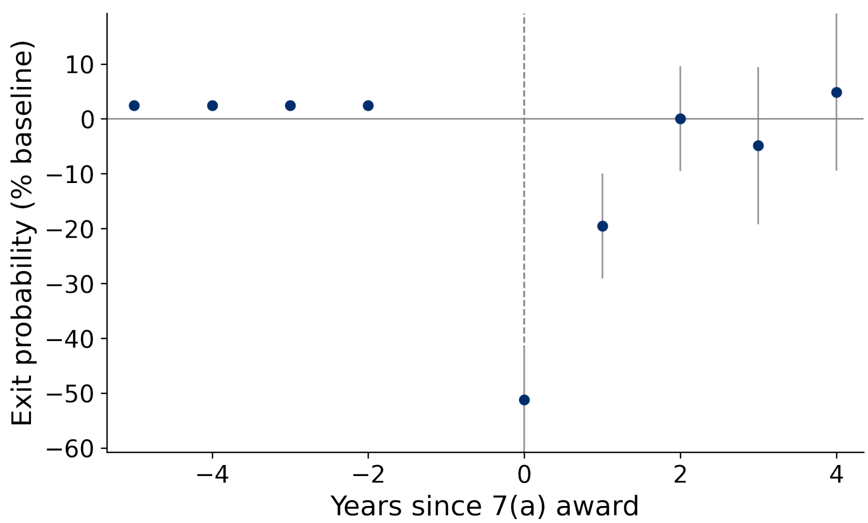

Note: This figure plots the event study estimates of a firm's exit probability after receiving 7(a), where event time is in years. The y-axis measures percentage drop in the conditional probability of exit in t for firms surviving until $t-1$. Pre-trends are nearly flat since pre-treatment conditional exit probability is identically zero for all firms selected into analysis. The average baseline exit probability in the pre-event year is 0.04. 95% confidence intervals are drawn using standard errors that are clustered by city-industry. Estimates cleared for disclosure by U.S. Census under disclosure ID 12684.

Summarizing the average effect over the five-year post-treatment window, the difference-

[19] Precise estimates pending disclosure review.

in-differences estimates in Table (4) show that 7(a) recipients are 0.9 percentage points less likely to exit each year relative to matched controls—a 22% reduction from the 4.1% baseline exit rate. The result is robust when we include industry and city fixed effects.

Table 4: 7(a) and Exit

	Coefficient	Exit Rate
Treatment Effect	-0.009	
	(0.001)	
Control Group (Post)		0.041
Treated Group (Post)		0.032
N	282,000	

Note: Difference-in-differences estimate of 7(a) receipt on annual exit probability over a five-year post-treatment window. Treatment effect represents the average treatment effect on the treated (ATT) relative to matched control firms in markets without 7(a) lending in the same industry. Controls include lagged average employment, revenue, and calendar-year fixed effects and firm fixed effects. Exit rates calculated from post-treatment period means. Standard errors clustered at city-industry level. Estimates cleared for disclosure by U.S. Census under disclosure ID 12684.

A key concern for our identification strategy is selection on unobservables that jointly predict 7(a) receipt and firm survival. Our design already conditions on a rich set of firm-level observables—revenue levels and growth, employment, and industry-year fixed effects—which capture the principal inputs and outputs of service-sector production. Remaining unobservables are likely to operate through financial channels. By program design, 7(a) recipients are firms that face greater difficulty obtaining conventional credit due to limited collateral or higher leverage (Office of the Comptroller of the Currency 2014b). This implies that untreated firms of comparable size and age within the same industry are financially more secure on average—less leveraged, better capitalized, and with more personal or relationship collateral. Consequently, absent program support, treated firms would be expected to have higher counterfactual exit risk than the average control group. Assuming that this pattern of selection is correct, our estimated effects present conservative bounds on the true impact of 7(a) guarantees on firm survival.

We conduct analogous analyses for the hotel industry, our primary case study. For exits, we estimate the same matched difference-in-differences specification, replicating the matching procedure across award years and firm-age bins. Control firms are hotels located in cities without any 7(a) lending during the sample period, ensuring no within-market exposure. Corresponding estimates will be reported as additional results following disclosure clearance from the U.S. Census Bureau.

4.2 Entry

Estimating entry effects requires a different approach, as potential entrants cannot be observed before they enter. We therefore shift to market-level analysis and exploit variation in local 7(a) credit supply using a shift-share instrument.

Specifically, we estimate the effect of 7(a) program lending volume on new firm entry using the following baseline specification:

$$y_{mt} = \beta^{entry} x_{mt} + \mathbf{W}'_{mt}\Gamma + \gamma_t + \mu_m + \varepsilon_{mt}, \tag{2}$$

where y_{mt} is the log number of new firm entries in market m and year t, x_{mt} is the log volume of 7(a) loans in dollars, γ_t and μ_m are year and city fixed effects and controls $\mathbf{W}_{mt}$ are average wage and total employment in the previous year. The coefficient of interest, β^{entry}, captures the elasticity of firm entry with respect to local 7(a) lending.

A simple OLS regression of y_{mt} on x_{mt} would be biased if local 7(a) lending responds to unobserved demand conditions that also influence firm entry. To address this endogeneity, we instrument local 7(a) loan volume with a shift–share measure that captures quasi-random variation in local exposure to national lending shocks. Our identification strategy follows Greenstone, Mas, and Nguyen (2020) in exploiting the highly localized nature of bank 7(a) activity: banks lend disproportionately in markets where they have established presence. When a bank experiences a transitory national shock—such as deposit growth—its local lending rises more in markets where it was previously active. We operationalize this by interacting each bank's pre-period market share with its subsequent national 7(a) lending fluctuations, excluding the market in question (leave-one-out):

$$Z_{mt} = \sum_b s^{\text{pre}}_{mb} \cdot \tilde{B}_{bt} \tag{3}$$

Here, s^{pre}_{mb} is bank b's share of 7(a) lending in market m in 2009, and $\tilde{B}_{bt}$ is bank b's residualized national lending in year t—demeaned and residualized against bank-specific trends and lags to isolate transitory, bank-specific supply shifts. To illustrate: TD Bank had sizable 7(a) presence in Philadelphia in 2009. When TD expands nationally (e.g., due to deposit growth), that expansion flows disproportionately to Philadelphia. If national bank shocks are uncorrelated with local conditions, this generates quasi-random variation in local credit supply.[20]

Table 5: 7(a) and Entry

	OLS	IV		
	(1)	(2)	(3)	(4)
log(SBA)	0.585	0.454	0.226	0.176
	(0.013)	(0.041)	(0.076)	(0.074)
Fixed effects	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Quasi-differencing			Yes	Yes
City trends				Yes
N	7,200	7,200	7,200	7,200

Note: This table presents our entry estimates, outcome variable is log(#entrants). Column (1) shows the OLS relationship, which likely reflects both supply and demand factors. Columns (2)-(4) present IV specifications with progressively richer controls. Our preferred specification (Column 4) includes city fixed effects, lagged average wage and employment controls, quasi-differenced values of IV, and city-specific trends to absorb local growth trajectories. Estimates cleared for disclosure by U.S. Census under disclosure ID 12684.

Our preferred specification includes city fixed effects and city-specific linear trends to absorb local growth trajectories. We also implement a quasi-difference specification to mitigate serial correlation in the instrument that might stems from national lending shocks. The instrumental variables estimates imply an entry elasticity of 0.176: a 10% increase in local 7(a) lending is associated with a 1.76% increase in new business formation. The IV estimates are substantially smaller than the OLS coefficient (0.585), consistent with positive selection—markets with more entry also attract more 7(a) lending.

We conduct analogous analysis for the hotel industry, our primary case study. For entry, the sparsity of hotel formation motivates a Poisson specification rather than log-linear IV; we estimate a Poisson pseudo-maximum likelihood model with the same shift-share instrument.

Together, these results suggest that SBA 7(a) loans meaningfully support the survival of existing firms and spur firm entry. We now turn to a structural model to evaluate the equilibrium and welfare implications of these changes.

[20] The variation in $\tilde{B}_{bt}$ stems from multiple sources, including a statutory increase in the program cap from \$18.75B to \$23.5B in 2015. While this policy change generates substantial variation, we do not rely exclusively on this single shock; our identification draws on the full distribution of idiosyncratic bank-year fluctuations over 2010-2019.

5 Model

We develop a model of hotels with external financing. This enables us to study how credit guarantees shape firm entry and exit, impacting prices and welfare.

The model unfolds in discrete time over an infinite horizon. It features two markets: a debt market, where hotels demand debt; and an accommodation market, where hotels compete for consumer demand.

The accommodation and debt market interact closely. In the accommodation market, hotels compete for consumers, which generates their profit streams. These profit streams are the key primitive in the debt market, where hotels obtain loans against their future profit stream. The availability of loans determines which new hotels can enter, and which existing hotels continue operating. This entry and exit in turn reshapes the structure of accommodation markets, impacting competition, prices, and profits.

We start with the hotel's profit stream and its core components in Section 5.1, while abstracting away from the details of the accommodation market. This stream forms the basis for lending decisions in the debt market.

Then, we describe the supply side of the debt market in Section 5.2. Lenders are identical and operate under perfect competition. They provide two types of loans: real estate loans, which are risk-free and modeled as proportional deductions from mean profits; and working-capital loans, which are risky since their recovery in default is limited. For working-capital loans, interest rates and maximal debt amounts, called debt limits, are determined endogenously within the model.

Subsequently, we present hotels as the demand side of the loan market in Section 5.3. Hotels take on debt to cover entry costs and adjust their borrowing over time based on realized profits. If a hotel's debt needs exceed its debt limit (determined endogenously), it defaults and exits the market. Conversely, lower borrowing costs and higher debt limits extend firm survival, raise profitability, and encourage new entry.

Finally, in Section 5.4 we microfound the profit stream by modeling competition among hotels in local accommodation markets. Market structure in this environment determines price competition, revenues, and the resulting profit stream that feeds back into debt dynamics.

We now turn to the components of the model in detail, beginning with the hotel's profit stream—the key primitive that links the accommodation market to the debt market.

5.1 Hotels as a Profit Stream

Hotel profits arise from accommodation sales and are subject to uncertainty, which affects hotels' loan terms.

In period t, a hotel's profit comprises an expected component with mortgage deduction, $\bar{\pi}_{jt}(1-a)$, and an innovation $\tilde{\pi}_{jt}$.

$$\pi_{jt} = \bar{\pi}_{jt}(1-a) + \underbrace{\pi^{\xi}_{jt} - \pi^{C}_{jt} + \pi^{\phi}_{jt}}_{\tilde{\pi}_{jt}}$$

$\bar{\pi}_{jt}$ is the hotel's ex-ante expected profit. This is determined by its characteristics (e.g., location, quality). Time dependence in expected profit arises from possible changes in the set of competitors.

The parameter a is the share of the firm's average profit that it pays for its mortgage. Within our model, real estate debt is assumed to be fully recoverable. It poses no risk to the lender and therefore a borrower's current state does not influence its terms.[21] We assume real estate debt is repaid using steady mortgage payments that are proportional to it's mean profit, based on an exogenous parameter a. This keeps our model tractable while still capturing scale: larger hotels, that earn higher profits, pay more in mortgages. The stability and long term nature of mortgage payments are empirically consistent with commercial real estate financing. See Appendix 9.7 for a detailed discussion of the two-part debt setup.

The stochastic component $\tilde{\pi}_{jt}$ represents various sources of uncertainty that drive variation in firm profits. The first term π^{ξ}_{jt} arises from idiosyncratic demand shocks in the accommodation market. The second term π^{C}_{jt} captures i.i.d cost shocks that come from operational hazards. These shocks explain excess exit among small firms.

The final term π^{ϕ}_{jt} reflects persistent profit differences due to unobserved managerial quality ϕ, which also affect hotel demand. Lenders learn ϕ gradually over time. This mechanism mirrors passive learning as in (Jovanovic 1982), where fixed types are revealed through observed outcomes. We assume that the distributions $D^{\pi^{\xi}}_{jt}, D^{\pi^{C}}_{jt}, D^{\pi^{\phi}}_{jt}$ of each of the profit shocks $\pi^{\xi}_{jt}, \pi^{C}_{jt}, \pi^{\phi}_{jt}$, respectively, are common knowledge to all lenders and firms.

Using these profit streams as primitives, we now describe the lending model and the

[21] As a result, changes in the borrower's state don't alter the price of real estate debt. Interest rate in this case would just be the risk free rate that compensates the lender exactly for the cost of raising funds. With the cost $(1+m)$ per dollar in the next period, the interest rate would be exactly m

key frictions it embeds.

5.2 Debt Market

Lenders provide loans backed by a hotel's future income stream. The debt market determines the interest rates and borrowing limits that shape the firm's entry, survival, and exit.

This part of the model links hotels' profit primitives—mean profitability and volatility—to credit allocation. This allocation also depends critically on the recovery rate in default. A higher recovery rate expands credit supply and relaxes borrowing limits, while limited recovery tightens allocation. Through this channel, the SBA 7(a) program affects both firm dynamics and market outcomes by improving lenders' recovery prospects.

The remainder of this section describes the credit market in two parts. Section 5.2.1 outlines the environment that determines interest rates and debt limits, while Section 5.3 presents hotels as borrowers whose entry and exit follow from these credit conditions.

5.2.1 Lending Environment

We now describe our empirical model of lending, starting with the lending environment.

Lenders allocate credit and set interest rates and debt limits, both shaped by the borrower's default risk. Throughout this discussion, we assume that lenders are identical, operate under perfect competition, and face frictionless refinancing, so lender identities do not matter.

Following standard corporate-finance and macro-credit models (e.g., Eaton and Gersovitz (1981) Hennessy and Whited (2005); Jermann and Quadrini (2012)), we model debt through a sequence of one-period contracts. This provides a tractable framework to capture two key ingredients of credit frictions in our environment: (i) the non-contingency of debt payments (ii) limited recovery in default. Later, we describe how these loans are a good representation of financing in our setting.

When a firm obtains a one-period loan, it repays the matured debt—principal and interest—at the end of the next period. It can repay this loan with profit alone, or it may obtain an additional loan to roll over the existing debt. Concretely, let L_{jt} denote a loan obtained by a borrower in period t. It is due for repayment in $t+1$ with added interest r_{jt}, after firm draws a profit π_{t+1} from its profit stream. The firm is said to roll over its debt into the future when it acquires new debt L_{jt+1} in $t+1$ to pay off its obligation

$L_{jt}(1+r_{jt})$:

$$L_{jt}(1+r_{jt}) = \begin{cases} \pi_{jt+1} + L_{jt+1} & \text{if } \pi_{jt+1} < L_{jt}(1+r_{jt}) \\ \min\{\pi_{jt+1}, L_{jt}(1+r_{jt})\} & \text{if } \pi_{jt+1} \geq L_{jt} \end{cases} \tag{4}$$

Thus, one-period contracts can be used to form long term repayment schemes. For example, consider a hotel that requires $1 million in advance cash flow. A lender offers debt at 10% interest through two one-period loans. Starting from a $1 million contract, the firm repays $1.1 million at the end of the first period, using $0.6 million in profit and a second loan of $0.5 million. It then repays $0.55 million at the end of the second period to settle the remaining debt.

Rollover in this environment is limited by an endogenous borrowing limit $L_{jt}^{\max}$. If a firm j in period t needs a larger loan than its limit $L_{jt}^{\max}$, it defaults. In default, the lender only recovers a fraction of the standing loan : ωL_{jt}. This limited recovery is the source of risk in debt contracts. Each one-period loan is priced with an interest rate that reflects its default risk, while the debt limit defines the maximum feasible borrowing consistent with lender participation.

This one-period structure captures the main empirical feature of working capital loans in service industries—revenues, as well as working-capital needs, arise regularly. Loan repayment closely follows the steady stream of firm income (Geelen 2024; Graham and Harvey 2001; Sufi 2009)[22]. At the same time, our framework necessarily abstracts away from banking relationships. Such relationships can create long-term rents for lenders and expand credit access for small firms (Chodorow-Reich 2014; Petersen and Rajan 1994). While relationships may expand credit supply, they don't fully mitigate default risk. In practice, small firms see large rates of default. Forbearance—temporary payment relief—is rare among most loans, which limits the extent to which relationships help in times of distress (Mourad, Schiozer, and Santos 2020), and distressed small firms typically default quickly[23].

We argue that even in the presence of relationship lending, small firms remain credit constrained due to default risk.[24] The 7(a) program is designed to relax these constraints

[22]In contrast, industries with lumpy or back-loaded payoffs—such as manufacturing or infrastructure—use long-maturity debt to align repayment with project returns (Geelen 2016; Stohs and Mauer 1996).

[23]Moreover, refinancing itself is a common feature among business loans, which dilutes the value of relationships, (Roberts 2015)

[24]Recent work highlights that the prospect of default, as opposed to continuation, can also reduce a firm's *inherent* value—for example, by lowering worker or managerial talent—so the costs extend beyond immediate recovery losses.

by improving lender recovery. Our model isolates these key forces—repayment rigidity and imperfect recovery—while maintaining a simple, tractable structure.

5.2.2 Interest Rates and Debt Limits

We now describe how the model determines (i) the interest rate r_{jt} and (ii) the equilibrium debt limit $L_{jt}^{\max}$.

Lenders set interest rates through a zero-profit condition, so rates reflect borrowers' default risk. The debt limit is the maximum amount of debt on which the lender breaks even. Both are equilibrium outcomes, but we present them sequentially for clarity.

A lender's information set allows it to assess the borrower's probability of default. We assume the lender observes the firm's expected next-period profit, its profit history, and its future credit limit. The profit history in t, defined as the set of profits until t : $\mathcal{H}_t := \{\pi_0, \cdots, \pi_t\}$, allows lenders to update their beliefs on the firm's type.[25] The lender also knows how debt limits L_{jt+1}^{max} would be determined in equilibrium. Given this information set, the lender chooses an interest rate.

Interest Rates Interest rates compensate the lender for the risk of borrower's default next period. Lenders are perfectly competitive, so this interest rate is determined by a zero profit condition. For now, we'll assume that debt limits $\{L_{jt}\}$ are already determined in equilibrium. We describe how they are determined in the next section.

Default in $t+1$ occurs if the borrower's required debt exceeds its debt limit. The borrower's next period liability is the amount it needs to borrow to pay off its debt and clear its costs:

$$L_{jt+1} = \max\{(1 + r_{jt})L_{jt} - \pi_{jt+1}, 0\}$$

L_{jt+1} increases with the size of current debt L_{jt} and the interest rate burden r_{jt} . If the borrower draws a sufficiently low profit in $t+1$, this liability will exceed it's debt limit L_{jt+1}^{max}. The next period's default probability, then, is the likelihood of low profit realizations that lead to default:

$$p_{jt+1}(L_{jt}, r_{jt}, \bar{\pi}_{jt+1}, \mathcal{H}_{jt}) = \Pr_{\pi_{jt+1}}\left(\pi_{jt+1} \leq (1 + r_{jt})L_{jt} - L_{jt+1}^{\max} \,\middle|\, \mathcal{H}_{jt}, \bar{\pi}_{jt+1}\right)$$

Given this default probability, the lender's payoff in the next period, π_{jt+1}^{lender}, is a lottery. It recovers the full loan, with interest, if the firm doesn't default. Otherwise, it only

[25] Longer histories lead to sharper inference, which improves credit terms for older firms. An entrant firm's history is empty, and the lender uses a public prior for inference on the firm's type. See Appendix 9.10.1 for details on Bayesian updating using profits.

recovers ωL_{jt}:

$$\pi^{lender}_{jt+1}(L_{jt}, r_{jt}, \bar{\pi}_{jt+1}, \mathcal{H}_{jt}) := \begin{cases} (1+r_{jt})L_{jt}, & \text{with probability } 1 - p_{jt+1}(L_{jt}, r_{jt}, \bar{\pi}_{jt+1}, \mathcal{H}_{jt}), \\ \omega L_{jt} & \text{with probability } p_{jt+1}(L_{jt}, r_{jt}, \bar{\pi}_{jt+1}, \mathcal{H}_{jt}) \end{cases}$$

where $\omega \in [0, 1]$ denotes the recovery rate in default.

In this setup, lenders face a tradeoff: higher interest rates increase payoffs upon successful repayment, but also raises the borrower's default probability p_{jt}.

Given perfect competition, the lender chooses the interest rate r_{jt} to earn zero profits in expectation. Let m denote the required return on funds—reflecting lenders' marginal cost of raising funds, or an opportunity cost. Then the interest rate satisfies:

$$\mathbb{E}[\pi^{\text{lender}}_{jt+1}(L_{jt}, r_{jt}, \bar{\pi}_{jt+1}, \mathcal{H}_{jt})] = (1+m)L_{jt}$$

In equilibrium, the interest rates r_{jt} rises monotonically with loan size L_{jt}. To see why: as L_{jt} increases, so does the default probability p_{jt} for a fixed interest rate, violating the break even condition. To maintain zero profit, lenders must charge higher rates to compensate for this increased risk. This further increases the default probability. However, for sufficiently low amounts of debt, the increase in maturity payoff more than offsets increased default risk from higher interest rates.

Beyond some threshold, no feasible interest rate satisfies the zero-profit condition for a borrower's debt. This places an upper bound on the firm's borrowing capacity.

Debt Limits When a loan becomes sufficiently large relative to a borrower's expected profits, it becomes infeasible. This is because a borrower that carries such a large debt defaults in most states. No interest rate compensates the lender enough to offset default risk without further increasing default probability. We define firm j's period t debt limit $L^{\max}_{jt}$ as the largest loan that is consistent with the zero-profit condition:

$$L^{\max}_{jt} \;=\; sup_{L_{jt}}\{L_{jt} : \mathbb{E}[\pi^{\text{lender}}_{jt+1}(L_{jt}, r_{jt}, \bar{\pi}_{jt+1}, \mathcal{H}_{jt})] \geq (1+m)L_{jt} \;\text{ for some } r_{jt}\} \tag{5}$$

$L^{\max}_{jt}$ depends on the debt limit in L^{max}_{jt+1} since the next period debt limit determines the firm's default probability. Additionally, it depends on the firms' mean profit next period, as well as its history, which allows the lender to infer default probability.

We show that under mild assumptions on the tails of the profit distribution, this limit exists and is unique (see Appendix 9.8.1). The intuition is that as loan size increases, repayment becomes increasingly reliant on extreme profit realizations. The thin-tail condition ensures such high-profit outcomes are sufficiently rare that no interest rate can offset the risk.[26]

Firms with higher expected profit $\bar{\pi}_{jt}$ have higher debt limits. These endogenous debt limits mirrors loan-to-value constraints commonly used in credit models (e.g., Kiyotaki and Moore (1997); Geanakoplos (1997); Buera, Moll, and Shleifer (2013)). Similarly, higher recovery rate ω , increases debt limit since better collateral reduces the lender's losses.

In equilibrium, the debt limit is pinned down as the value of debt at which the profit maximizing interest rate exactly satisfies the zero profit condition:

$$\mathbb{E}[\pi^{\text{lender}}_{jt+1}(\bar{\pi}_{jt+1}, \mathcal{H}_{jt}, L_{jt}, r^*_{jt})] - (1+m)L_{jt}\big|_{\ L_{jt}=L^{max}_{jt}} = 0$$

where r^*_{jt} satisfies the interior first order condition to be the profit maximizing interest rate:

$$\frac{d}{dr_{jt}}\mathbb{E}[\pi^{\text{lender}}_{jt+1}(\bar{\pi}_{jt+1}, \mathcal{H}_{jt}, L_{jt}, r_{jt})]\big|_{\ L_{jt}=L^{max}_{jt}, r_{jt}=r^*_{jt}} = 0$$

As long as the firm faces a finite debt limit asymptotically, the equilibrium system has a finite debt limit in every period. In our setting, the lenders' beliefs about a borrowers' type converge to a stationary distribution around its true value. In this stationary environment, we show that a finite, unique debt limit exists (see appendix 9.8.1).

5.2.3 7(a) and Loans

The 7(a) program operates by reducing the lender's losses when a guaranteed borrower defaults. In our model, this is captured as an increase in the effective recovery rate ω for guaranteed loans, which alters loan pricing and debt limits in equilibrium.

First, higher recovery rates imply higher debt limits (See Appendix 9.8.2). Intuitively, higher recovery rates reduce lenders' expected losses in default, allowing them to extend larger loans while still breaking even.

Second, reduced default losses allow lenders to set lower interest rates. Figure 4 illustrates both effects: the guarantee shifts the (profit) threshold for firm exit by increasing debt limit and decreasing interest rates. These responses imply that guarantees pass through to borrowers as increased debt access and lower borrowing costs.

We assume the required cost m is unaffected by 7(a), consistent with the program's

[26] Normal and exponential distributions satisfy this criterion, as do many standard fat-tailed distributions such as the Pareto.

Figure 4: Effect of 7(a) Guarantees on Credit Constraints

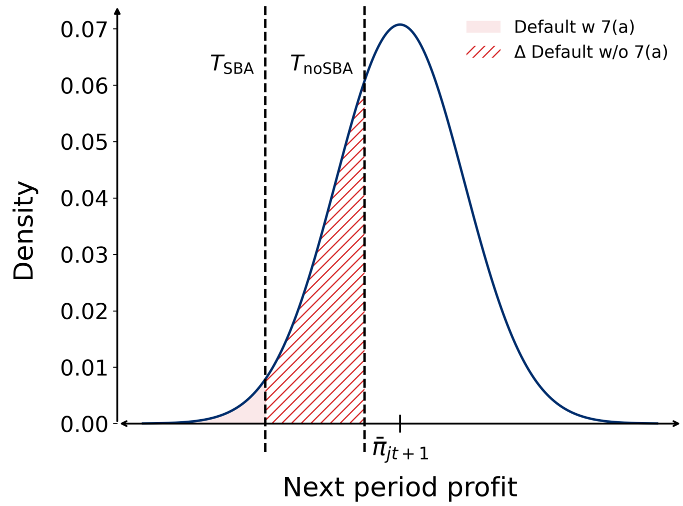

Note: Model-implied exit probability for the marginal non-entrant under estimated parameters. The y-axis shows the probability density of next-period profit; regions left of the vertical line represent profit threshold for default without and with 7(a). The 7(a) guarantee increases the credit limit and decreases interest rates. The threshold profit $T := (1 + r_{jt})L_{jt} - L_{jt+1}^{max}$ defines the profit realizations below which the firm defaults.

small scale relative to total credit supply.[27]

Having characterized the credit supply side, we now turn to the hotels that demand debt, which influences their value, entry and exit.

5.3 Borrowers

Hotels' ability to finance dictates their profit, entry and exit. They acquire debt to pay for costs, and debt evolves according to the terms determined in the debt market.

Upon entry, hotels need debt to pay for a fixed cost. Once active, they earn a profit stream and use it to pay off their initial debt. They reduce debt if profits are high, but when profits are low relative to costs, they demand new debt. As described before, when debt requirements exceed debt limits, they default.

We now concretely describe how hotels transition between solvency and default, then define their value function and entry rule.

[27]Stillerman (2024) explores complementary lender-side responses, specifically moral hazard of banks.

5.3.1 Dividends and Exit

A hotel starts period t with liability $L_{jt-1}(1+r_{jt-1})$ from past operations. After profit π_{jt} is realized, the hotel must repay this liability in full. Hotels cannot raise equity and rely solely on debt to finance their costs— negative profits and/or prior debt.[28] Each period, they must clear their accounts using current profits or new borrowing. This structure reflects the inflexible working-capital demands of small-business hotels, where expenses must be met continuously to remain in operation. (American Hotel & Lodging Association 2020, Federal Reserve Banks 2019).

After servicing debt, remaining profits, if any, are disbursed as dividends:

$$D_{jt} = \max\{\pi_{jt} - L_{jt-1}(1 + r_{jt-1}), 0\}$$

If profits fall short, the hotel demands new debt $L_{jt} = \max\{0, L_{jt-1}(1 + r_{jt-1}) - \pi_{jt}\}$ to pay for costs and previous debt. Hotel defaults when the required borrowing exceeds the equilibrium debt limit: $L_{jt} > L_{jt}^{\max}$. We treat default and exit as equivalent— firm shutdown occurs only through financial distress. Upon default, the hotel loses ωL_{jt}.

Our assumption of debt and exit being equivalent events is strong. If some firms exit for reasons unrelated to financial distress, our framework will overstate the role of credit frictions in explaining exit. Empirically, credit delinquency strongly predicts small-business exit and is the primary cause for exit. (Metaxas and Romanopoulos 2023, Gu and Gao 2000). SBA policy reinforces this: lenders may demand the guarantee from SBA if the borrower defaults on an installment payment for more than 60 days. This triggers recovery proceedings that effectively force exit. Additionally, most small business loans lack softer covenants that can transfer some control rights to lenders rather than forcing default.

Having described the hotel's movement across states of debt and their payoffs, we now define the hotel's value.

5.3.2 Firm value

A hotel's value determine its entry, as described in the next section. It's value is simply the discounted sum of its payoff stream. Payoffs are positive when profits generate

[28] Most small U.S. hotels operate as independent local businesses that lack access to public or institutional equity markets. Consistent with broader small-business evidence, they rely almost entirely on bank and SBA-guaranteed loans to finance fixed and working-capital needs. Venture and private-equity participation remains negligible, and the majority of firms carry outstanding debt rather than external equity funding (U.S. Department of the Treasury 2023; U.S. Small Business Administration 2024b).

dividends but negative upon default when the firm pays loan recovery costs. We write value recursively:

$$V(\pi_{jt}, L_{jt-1}(1+r_{jt-1}), \mathbf{S}_{jt}) = \begin{cases} D_{jt} + \beta\,\mathbb{E}[V(\pi_{jt+1}, L_{jt}(1+r_{jt}), \mathbf{S}_{jt+1})], & \text{if } L_{jt} \leq L_{jt}^{\max} \\ -\omega L_{jt-1}, & \text{otherwise.} \end{cases} \tag{6}$$

The additional terms in the state vector state vector $\mathbf{S}_{jt} := \{\bar{\pi}_{jt+1}, \mathcal{H}_{jt}, n_{jt}, \psi(n_{jt})\}$ include history $\mathcal{H}_{jt}$ and expected profit for next period $\bar{\pi}_{jt+1}$, which determine the cost of debt. Additionally the firm keeps track of the number of firms n_{jt} and a market structure evolution function $\psi(n_{jt})$.

Market structure evolution matters because changes in the set of competitors affects future continuation profits. In principle, a hotel's continuation value depends on the full distribution of competitors' financial positions, which determines exit, and a rational expectation on entry decisions. Since tracking every firm's financial position is infeasible, hotels form expectations using aggregate statistics. Following the oblivious equilibrium approach (Weintraub, Benkard, and Van Roy 2008) we summarize market evolution with a function that maps number of firms to entry/ exit probability. $\psi(n_{jt}) = (P_e, P_x|n_{jt})$ denotes the probability of exit (P_x) and entry (P_e) for hotel j's market that has n_{jt} firms.[29] Hotels treat the current number of firms n_{jt} and the equilibrium market evolution function $\psi(.)$ as sufficient statistics to forecast how market structure evolves and affects future profits.[30]

5.3.3 Entry

Potential entrants enter a market when their expected value exceeds an equity requirement. One potential entrant arrive each period in every market, with its amenity value drawn from distribution $\mathcal{D}_{\text{entrant}}$ that determines its profit. It decides whether to pay for an entry cost $F(.)$ that depends on the firm's expected profit to capture scale in entry costs. Consistent with small-business loan structure, entrants finance a fraction of their entry cost $sF(.)$ with own equity and $(1-s)F(.)$ with debt (Office of the Comptroller of the Currency 2020). The potential entrant faces an error in their anticipated firm profit

[29] We'll describe the entry protocol as "announce in t to enter in $t+1$" so, P_e here is the probability of announcement. Additionally, we'll hold this probability fixed across firm sizes for now, for computational tractability in counterfactuals

[30] We assume that firm's use an "average" percentage drop in profits going from a set of n to $n+1$ firms to calculate change in profits from entry or exit, where we take averages implied by the estimated demand model (later) and the observed set of firms in the data. That is, when we estimate the model, we use the average profit change from actual set of firms seen in the support. In counterfactuals, we assume these and check for consistency with the generated one. See section on equilibrium 5.5

of $\eta_j \sim D_\eta$. This error term explains entry by low profit firms through the model. We assume entrants announce their decision before period-t lending concludes, which lets us separate credit decisions from immediate market structure changes.

An entrant with expected ex-ante profits $\bar{\pi}_{jt}$ at the time of entry, and idiosyncratic entry shock η_j enters when expected firm value exceeds the equity requirement:

$$\mathbb{E}\left[V(\bar{\pi}_{jt} + \eta_j, (1-s)F_{jt}(1+r_{jt}), \mathbf{S}_{jt})\right] - sF_{jt} > 0 \tag{7}$$

where $F_{jt} = F(\bar{\pi}_{jt} + \eta_j)$ is the total (anticipated) entry cost. $(1-s)$ of this cost is paid using debt, and $r_{jt} = r(\bar{\pi}_{jt} + \eta_j, (1-s)F_{jt}, \mathcal{H}_{jt})$ represents the interest on this anticipated debt burden. The vector $\mathbf{S}_{jt} = \{\bar{\pi}_{jt+1} + \eta_j, \mathcal{H}_{jt}, n_{jt}, \psi_{jt}\}$ describes the remaining state variables, as in 5.3.2. After a firm enters, its actual loan and profits are free of this error η_j.

Since credit terms affect entry incentives, programs like 7(a) can expand entry by improving credit access for beneficiaries.

5.3.4 Borrowers and 7(a)

The 7(a) program's improves firms' loan terms through higher debt limits and lower interest rates. This allow beneficiaries to accrue less debt, roll it over for longer, default less often, and earn higher long-run profits. These improvements raise firm value and increase the likelihood of entry.

These advantages also affect non-recipients through increased competition. Whether the program generates net social welfare gains depends on whether beneficiaries expand total market demand or simply reallocate existing business. This distinction hinges on demand elasticity and the intensity of price competition, which we now specify through the accommodation market structure.

5.4 Accommodation Market

The accommodation market generates hotel profits. This creates the profit streams that working capital loans are provided against. In turn, the lending market dictates hotel entry and exit, which implies the set of hotels that compete for accommodation services.

Every period, consumers arrive in a location and seek accommodation. Hotels in the location compete on prices. In equilibrium, this determines firms' profits and consumer surplus.

Concretely, a consumer i derives the following utility from choosing hotel j in nest $\mathcal{J}(j)$ at time t:

$$u_{ijt} = \delta_j + \phi_j - \beta p_{jt} + \xi_{jt} + \nu_{\mathcal{J}(j)t} + (1-\rho)\epsilon_{ijt}. \tag{8}$$

Here, ϵ_{ijt} are i.i.d. type-1 extreme value shocks, and $\nu_{\mathcal{J}(j)t}$ is drawn from the Cardell distribution (Cardell 1997). β measures disutility from price, ρ governs within-nest correlation of shocks (Miller and Weinberg 2017) . δ_j represents the amenity value of hotel j, reflecting quality, location, brand, etc., ϕ_j is the firm-specific managerial quality, and ξ_{jt} is an i.i.d. demand shock.

These preferences give rise to a nested logit demand system, where consumers substitute more readily within categories (motels) than across categories (motels vs. hotels). The nest parameter ρ controls within-category substitution: higher ρ means new hotels primarily steal demand from similar competitors, while lower ρ allows entrants to expand market size by attracting marginal consumers. Our model features 4 nests: hotels, motels, a catch-all nest "other" for campgrounds, B&Bs etc, and an outside option in a separate nest. This classification captures meaningful substitution patterns across accommodations of different types. Details of data and market construction are provided in Appendix 9.3.

For a given set of prices, we get the following product shares s_{jt}:

$$s_{jt} = \frac{\exp\left(\frac{v_{jt}}{1-\rho}\right) D_{gt}^{-\rho}}{1 + \sum_g D_{gt}^{1-\rho}}, \quad D_{gt} = \sum_{k \in g} \exp\left(\frac{v_{kt}}{1-\rho}\right), \quad v_{jt} = \delta_j + \phi_j - \beta p_{jt} + \xi_{jt}.$$

where g is shorthand for product j's nest $\mathcal{J}(j)$.

We assume that hotels observe all non-price utility components for all k firms in the market ($\{\delta_k, \xi_{kt}, \phi_k\}_k$) before setting prices. Each hotel sets p_{jt} to maximize profits in a Nash-Bertrand game:

$$p_{jt}^* = \arg\max_{p_{jt}} \quad s(p_{jt}|\mathbf{P}_{-jt})(p_{jt} - c_{jt}),$$

where $s(\cdot)$ denotes the market share function, suppressing dependence on non-price variables, and c_{jt} is marginal cost, $\mathbf{P}_{-jt}$ the vector of other firms' profits.

This yields profits based on the market size M_{jt} . We define variable profits Π_{jt} as $\Pi_{jt} = (p_{jt} - c_{jt}) \cdot s_{jt} \cdot M_{jt}$.

Now, we use a linear approximation around the firms profit when its own shocks are zero. This results in additively separable terms for firm type and demand that are consistent with the lending model:

$$\Pi_{jt} \approx \bar{\pi}_{jt}\left[1+\kappa_{jt}(\phi_j+\xi_{jt})\right], \qquad \kappa_{jt} = \frac{\partial \ln s_{jt}}{\partial v_{jt}} \tag{9}$$

here, κ_{jt} is the semi-elasticity of j's demand with respect to its utility, $\bar{\pi}_{jt}$ is the profit evaluated in absence of all demand shocks, and own type deviation—$\bar{\pi}_{jt} \equiv \Pi_{jt}\big|_{\phi_j=0,\boldsymbol{\xi}=\mathbf{0}}$. The profit deviations scale approximately linearly with demand shocks and types $\pi^{\xi}_{j}t = \bar{\pi}_{jt}\cdot\kappa_{jt}\phi$, $\pi^{\phi}_{jt} = \bar{\pi}_{jt}\cdot\kappa_{jt}\phi$.

This approximation captures the first-order effects of a hotel's own demand shocks and firm type, while abstracting from higher-order interaction terms and dependence on other firms' shocks. The linearization provides a tractable, additively separable representation of profits that links directly to the credit market.

These profits, after subtracting mortgage payments and i.i.d cost shocks, form the profit stream. Profit streams are the foundation for lending decisions in our model: they determine the credit available to each hotel, the distribution of dividends, and the likelihood of entry or exit.

5.5 Equilibrium

We adopt a stationary, oblivious rational-expectations equilibrium. The key equilibrium requirements are that lenders and hotels have rational expectations about their beliefs. Recall that lenders hold beliefs about equilibrium (next-period) debt limits (see 5.2.2). Hotels hold beliefs about the equilibrium entry and exit probabilities given current number of firms (see 5.3.2) while calculating their firm value and deciding entry.

In equilibrium, lenders' and hotels' beliefs should be consistent with implied equilibrium outcomes. That is, lenders' beliefs about equilibrium debt limits should coincide with the actual debt limits that emerge from lenders' zero profit condition, given those beliefs. Hotels' beliefs about average entry and exit probability should be consistent with the (average) entry and exit probabilities generated by the model.

We describe these equilibrium conditions as fixed points of two operators: an operation that maps beliefs about debt limits into implied debt limits, and an operator that maps entry-exit probability beliefs into implied average probabilities.

5.5.1 Credit Limit Consistency

Let $L^b(\bar{\pi}, \mathcal{H})$ denote a function that denotes lenders' beliefs of equilibrium debt limits given arguments $(\bar{\pi}, \mathcal{H})$. Then, we define an operator T_{credit} which gives an implied debt limit function based on the lender's beliefs and their break even condition:

$$T_{\text{credit}}[L^b](\bar{\pi}_{jt+1}, \mathcal{H}_{jt}) = \sup_L \{L \geq 0 : \exists\ r \text{ s.t. } \mathbb{E}[\pi^{\text{lender}}(r, L, \bar{\pi}_{jt+1}, \mathcal{H}_{jt} \mid L^b(.))] = 0\}.$$

This is how lenders set debt limits when they believe that next period equilibrium debt limits are given by $L^b(\bar{\pi}_{jt+1}, \mathcal{H}_{jt})$.

Then, the equilibrium set of debt limit beliefs satisfies a fixed point of the above operator:

$$T_{\text{credit}}[L^{b^*}](\bar{\pi}_{jt+1}, \mathcal{H}_{jt+1}) = L^{b^*}(\bar{\pi}_{jt+1}, \mathcal{H}_t)$$

This condition ensures rational expectations: the equilibrium choices made by lenders are consistent with their beliefs.

5.5.2 Market Evolution Consistency

The second fixed point ensures consistency in hotels' beliefs. Let $\mathcal{G}(\psi(.))$ denote the operator that maps the equilibrium beliefs about the market evolution function $\psi^b(n_{jt})$ into implied average probabilities of entry and exit for markets with n_{jt} firms:

$$\mathcal{G}(\psi^b(n_{jt})) = \{\text{implied average transition probabilities when hotels act under } \psi^b(.)\}$$

Then, in equilibrium, the implied average market transition probabilities must be the same as the hotels' beliefs, i.e., the belief must solve a fixed point problem in the map:

$$\psi^{b^*}(n_{jt}) = \mathcal{G}(\psi^{b^*}(n_{jt})$$

Note that, this map depends on all decisions made in the accommodation market, that generates profit streams, and in the lending market, where credit allocation and the profit stream give rise to entry and exit. We summarize these decisions for completeness below. The appendix 9.9 describes the equilibrium conditions more formally.

(1) Accommodation market actions : Hotels compete in Nash–Bertrand equilibrium given the nested-logit demand system, yielding equilibrium prices, market shares, and operating profits (see 5.4).

(2) Debt accounting and rollover: Debt rolls over across periods according to realized profits and interest rates, when approved (see 4).

(3) Lenders' zero-profit condition: Lender interest rates satisfy the zero profit condition, that also gives rise to debt limits (see 5.2.2).

(4) Default and exit. Firms default and exit when their required rollover exceeds their equilibrium debt limit (see 5.2.2).

(5) Value recursion. Firms earn dividends from profits after debt accounting. Firm value satisfies the recursive Bellman equation (see 6).

(6) Entry. Potential entrants draw amenity value and entry shock (δ, η) and enter if their entry criterion in equation 7 is satisfied.

6 Identification and Estimation

We estimate the model in stages, drawing on three primary sources of variation: hotel revenues, hotel entry, and hotel exit. Additionally, we use loan outcomes to infer lender costs.

Before outlining the estimation steps, we summarize the model parameters to be estimated and their role in the model. Consumer preferences (β, ρ) and firm characteristics create variable profits. Volatility in demand and costs, and the dispersion in types, dictated by their distributions (D_{ξ}, D_c, D_{ϕ}), creates uncertainty in profits, while mortgage deductions a reduce mean profits. These profit streams imply credit limits and interest rates, which dictate exit. Firm entry depends on entry costs $F(.)$, the (amenity) distribution of entrants $D_{entrant}$ and the distribution of entry shocks D_{η}.

Now we lay out a road-map of the various stages of estimation:

Accommodation Market

1. Consumer demand: We estimate guest preferences in the accommodation market first in Section 6.1.1 using standard approaches.

2. Profits and demand shocks: Once we estimate the guests' demand model, we use it to calculate mean profits and their demand deviations in Section 6.1.2.

Debt Market

3. Lender Costs and Losses: We estimate these directly from data on 7(a) loan outcomes in Section 6.2.1.

4. Entry: Then, we estimate entry parameters using maximum likelihood on the observed entry patterns in Section 6.2.2.

5. Profit Primitives: Finally, we estimate the remaining components of the profit stream by using maximum likelihood on observed exit, in Section 6.2.3. At a high level, exit patterns identify features of the profit stream because our lending model maps the profit stream into (a distribution of) firm exit. We solve the model to explicitly create this map, and use it to recover the underlying primitives.

This sequential approach maintains computational tractability while allowing us to exploit model structure for precision. In particular, the maximum-likelihood estimation in the debt market leverages the model's equilibrium relationships to extract the most information from the data.

6.1 Accommodation Market

6.1.1 Consumer Demand

We estimate consumer preferences for accommodation demand using cross sectional variation in prices and product shares.

Recall that markets are city-year pairs, with the largest (66) metropolitan areas subdivided into smaller markets to ensure meaningful product differentiation. Following Farronato and Fradkin (2022), we set market sizes to two times the calculated quantity of sales.

Using the Berry inversion (Berry 1994), difference in (log) product shares and the (log) outside share yields consumer utility, along with a term for within-nest shares. Using our specified guest utility (see 8), the estimating equation for demand is:

$$\log s_{jt} - \log s_{0t} = X_{jt}\alpha \; - \; \beta\, p_{jt} \; + \; \rho \log s_{j|g,t} \; + (\delta_j + \phi_j - X_{jt}\alpha) + \; \xi_{jt}\,. \tag{10}$$

where s_{jt} is hotel j's share in market t, , $s_{j|g,t}$ represents j's share market within its nest g, s_{0t} is the share of the outside good, X_{jt} are observable characteristics—hotel average rating, number of stars, and a brand dummy— which capture a portion of the hotels' amenity value and type (δ_j, ϕ_j). p_{jt} is price and ξ_{jt} represents unobservable hotel quality. The key parameters of interest are (β, ρ).

Identification Identification for price sensitivity comes from correlation in prices and market shares, while correlation in within-nest shares and product level shares identifies the nest parameter. In our estimation, this variation is fully cross-sectional, since we only have prices from 2024, when we collected our price data.

The key identification challenge is price endogeneity: hotels with higher unobserved quality charge higher prices, creating correlation between ξ_{jt} and p_{jt}.

We address price endogeneity using differentiation-based instruments that use other competitors' characteristics to create measures of competitive pressure. This follows Berry, Levinsohn, and Pakes (1995), Bresnahan, Stern, and Trajtenberg (1997) and Gandhi and Houde (2019) who use functions of rival product characteristics to construct instruments.

The first set of instruments exploits variation in number of competitors around a hotel. Specifically, we measure the number of hotels j competing with i by lying within one standard deviation of the inter-hotel distance distribution: $Z^G_{jt} = \sum_{k \in \mathcal{J}_m \setminus \{j\}} \mathbb{1}\{d_G(j,k) \leq h_G\}$, where d_G is geographic distance and $h_G = \mathrm{sd}\big(\{d_G(k,\ell)\}_{k,\ell}\big)$ is the standard deviation of all pairwise inter-hotel distances pooled across markets, capturing typical geographic dispersions.

The second set of instruments measures competitive pressure by counting competitors within the market that are 'nearby' in attribute space : (i) the number of same-nest, same-star hotels within h_G, and (ii) the number of hotels within one standard deviation of a quality q_j index derived from review text (see Appendix (9.3.6)). Formally, $Z^q_{jt} = \sum_{k \neq j} \mathbb{1}\{|q_k - q_j| \leq h_q\}$, where $h_q = \mathrm{sd}(\{q_k\}_k)$ is the standard deviation of the quality index across all hotels, capturing typical quality dispersion in sample. This instrument tracks competing close substitutes in perceived quality.

To address endogeneity in nest shares, we include number of products in a nest, as well as city-level accommodation tax rates as instruments. These shift the value of the inside nest relative to the outside option.

The instruments are relevant as they capture how "crowded" a hotel's competitive environment is. First-stage F-statistics exceed 50, and over-identification tests fail to reject instrument validity at the 95% level.

The exclusion restriction requires

$$\mathbb{E}[\xi_{jt} \mid Z_{jt}] = 0 \tag{11}$$

We argue that our instruments satisfy exclusion as they induce price variation only from supply side competition. For a given hotel, its competitors' characteristics affect demand only through induced changes in prices, not directly through the utility provided by the hotel. All instruments are leave-one-out.

An additional challenge in current estimation is that revenue data comes from the 2022 Census Business Register[31], while our scraped prices are from 2024. The timing mismatch creates measurement error in prices, but our instrumental variables strategy remains valid under classical measurement errors (see 9.10.1). Plausible forms of non-classical measurement error would bias our price coefficient toward zero, making our estimates conservative.

We estimate the system of equations in (11) using a 2-stage GMM procedure, which improves takes the variance-covariance matrix of moments from a first stage estimate (2SLS) and uses them to construct optimal GMM weights for the second step of estimation.

Table 6: Demand Estimates: Logit vs. Nested Logit (OLS and IV)

	OLS–Logit	IV–Logit	OLS–Nested	IV–Nested
Price coefficient β	-0.003	-0.05	-0.000	-0.015
	(0.00)	(0.006)	(0.039)	(0.003)
Nest parameter ρ	—	—	0.755	0.656
			(0.03)	(0.130)
Instruments	No	Yes	No	Yes
Nests	No	No	Yes	Yes
N	7200	7200	7200	7200

Notes: Coefficients shown with standard errors in parentheses clustered at city level. β is the marginal utility of price; ρ is the within-nest correlation parameter ($0 \leq \rho < 1$). Estimates cleared for disclosure by U.S. Census under disclosure ID 12761.

Results Estimates are shown in Table (6). These imply high price sensitivity for accommodations. Median own-price elasticity across all hotels is -4.36 and the mean is -5.23. This is consistent with prior estimates in the literature that find elasticities in the range of -5 to -7 (Farronato and Fradkin 2022). The within-nest correlation parameter $\rho = 0.66$ implies strong substitution within accommodation types. The implied (average) market elasticity of -1.74 indicates that while consumers readily substitute among properties, overall lodging demand is moderately elastic— consistent with significant business stealing.

[31]This is the current latest version of revenues. Future iterations will use updated revenue data, obviating this challenge.

6.1.2 Recovering Profits and Demand led Deviations

Having estimated consumer preferences, we now translate observed revenues into firm-level variable profits. From these, we recover expected firm profits and demand-driven deviations, which are components of the underlying profit stream. Mean profits are used in later estimation stages.

Recovering Profits First we translate revenues into firms' variable profits.

A challenge for recovering profits is that we don't observe prices between 2010-2019 for small hotels. Prices are only available for 2024 through our data collection exercise. To overcome this, we predict marginal costs and use the Nash–Bertrand model of price competition to back out prices consistent with profit maximization. Concretely, given marginal cost and revenue for every firm, we can solve for the vector of prices that would be consistent with each firm maximizing their profit.

Two assumptions are required for this step. First, the predictive model of marginal costs must approximate firms' true costs during 2010–2019. Second, we assume that consumer preferences during this period are stable and identical to those estimated from 2022 data.

We predict marginal costs using local labor costs. Here, we want to capture how local labor markets drive the price of accommodations. Labor represents over 50% of variable hotel costs (U.S. Bureau of Economic Analysis 2025), and predicts average local prices and electricity costs, making wages a significant and natural cost predictor.

To build a predictive model, we back out marginal costs using the estimated demand system and supply model (with data from 2022 and 2024). Using these recovered costs as data, we regress marginal costs on local wages. We then use the estimated linear relationship to predict marginal costs for 2010–2019. This linear regression yields median predicted costs of $71 per room for motels and $94 for hotels through our sample (in 2019 dollars), which is consistent with benchmarks and our actual estimates.[32]

Next, we use these marginal costs to infer equilibrium prices. We exploit the first-order conditions for profit maximization: given demand parameters and marginal costs, we solve for the vector of prices that is consistent with Nash-Bertrand equilibrium. Let $\mathbf{s}(\mathbf{p})$ denote the $J \times 1$ vector of market shares and $\Delta(\mathbf{p})$ the Jacobian with elements $\Delta_{jk} =$

[32] These figures are somewhat elevated relative to rooms department costs alone (which industry data suggest averaged $35-50 for limited-service properties in this period). This is an artifact of the high marginal costs implied by our estimated lerner markups in 2022. This reflects that our recovered costs include all variable expenses and incorporate credit-constraint pricing wedges.

$\partial s_j / \partial p_k$. The Bertrand first-order conditions yield:

$$\mathbf{mc} = \mathbf{p} + \Delta(\mathbf{p})^{-1}\mathbf{s}(\mathbf{p}), \tag{12}$$

where we impose single-product ownership (fewer than 7% of markets have multi-product firms). We solve this J-equation system using Newton's method for each market-year pair. Details are provided in the Appendix (9.10.2).

With prices, revenues and marginal costs, we compute firm-level variable profits as: $\Pi_{jt} = q_{jt}(p_{jt} - mc_{jt})$.

To obtain mean profits, we make a simplifying assumption: each firm's expected ex-ante profit is constant over its lifetime. We use the firm's average profit as an estimate for this lifetime mean[33]:

$$\bar{\pi}_j := \frac{\sum_t \Pi_{jt}}{\sum_t \mathbb{1}\{j \text{ is active}\}}$$

This simplification abstracts from within-firm profit changes caused by evolving market structure. Such abstraction could bias parameters inferred from exit, which rely on variation across the profit distribution. However, we argue that most hotels' relative positions in the profit distribution remain stable over time. Hence, potential bias from ignoring within-firm movements in profitability is likely limited.

Results: This yields firm-level markups and mean profits averaging 34% of revenue (median 29%), which are broadly consistent with industry benchmarks of 25–30% gross operating margins (STR / CoStar Group 2024). Profit and revenue rankings align tightly within markets—fewer than 1% of firms differ by more than one rank position.

Recovering Demand led Deviations in Profit With the mean profits $\bar{\pi}_j$ in hand, we now estimate demand-driven profit deviations, π^{ξ}_{jt}.

Our linear approximation in equation (9) implies that demand shocks change profits proportional to the mean profit, scaled by the semi-elasticity of substitution. Empirically, we find that profit variance scales linearly with mean profit in a constant proportion. This is consistent with having similar semi-elasticities across markets.

We therefore specify demand shocks as proportional to mean profit:

$$\pi^{\xi}_{jt} = \bar{\pi}_j \cdot \nu^{\xi}_{jt}, \qquad \nu^{\xi}_{jt} \sim \mathcal{N}(0, \sigma^2_{\pi^{\xi}}),$$

This specification allows us to pool information across firms, improving efficiency. We

[33]Note that, we cannot differentiate between expected profit and firm type. The two remain together during estimation, and the econometrician can account for uncertainty in types.

estimate this by maximum likelihood. The resulting demand deviations feed into later stages of estimation.[34]

Results: Maximum-likelihood estimates yield $\sigma_{\pi^{\xi}} = 0.37$, implying that the innovation in profits have a standard deviation of 37% of mean —consistent with industry reports of high volatility in hotel profits.

6.2 Debt Market

So far, we have recovered mean profits and demand shocks from hotel revenue data. We now turn to hotel entry and exit data to estimate the remaining model parameters.

Estimation proceeds in three stages.. First, we estimate the recovery parameter and the opportunity cost (ω, m) directly from loan outcomes (Section 6.2.1). Next, we use variation in firm entry across their profits to recover the entry cost function $F(.)$ and other entry parameters— $D_{\eta}, D_{entrant}$. For tractability, we estimate the entry-cost function in two steps, as described in Section 6.2.2. Finally, we use exit variation to recover the remaining components of the hotels' profit stream—(a, D_{ϕ}, D_C)—in Section 6.2.3.

We specify parametric forms for each of these components, allowing estimation by maximum likelihood. This approach lets us exploit the full structure of the model for identification and precision.

6.2.1 Funding Costs and Recovery

We first estimate the parameters that characterize the lender's one-period problem. These parameters enter directly into the debt-market block of the model.

The cost of raising funds m, and recovery rate ω are recovered directly from the data on 7(a) loan outcomes. We infer the recovery rate from the distribution of charged-off losses in the data and set the value of ω equal to the empirical mean of that distribution, 0.61.

We estimate the cost of funds m using the lenders' zero-profit condition. Because we observe principals, interest rates, and realized losses, we can back out the average cost that implies zero expected profit on the portfolio of hotel loans. This yields an implied cost of funds of 1.09 per dollar lent.

[34] Note that, this captures changes due to market structure as well, which partially compensates for our earlier choice of keeping profit averages fixed. While running counterfactuals, however, we don't currently account for how this variance changes.

6.2.2 Entrant Distribution and Entry Threshold

Now, we estimate firm entry parameters in 5.3.3: the fixed cost function $F(.)$, distribution of entry shocks D_η, and the distribution of (amenity values of) underlying entrants $D_{entrant}$. We specify parametric distributions for these objects and estimate them by maximum likelihood.

As an estimation device, we first recover a profit threshold for entry, π^*, rather than estimating fixed costs directly. We then use this threshold to infer the fixed-cost parameters. This two-step procedure exploits the model's structure and allows us to divide estimation into smaller, tractable components. In this section and beyond, we use firms' empirical mean profit $\bar{\pi}_j$ for estimation, as discussed in 6.1.2. For notational consistency with the model, we retain a time subscript $\bar{\pi}_{jt}$, though estimation uses $\bar{\pi}_j$.

To capture scale effects in entry costs, we parameterize entry cost to be linear in mean profit: $F(\bar{\pi}_{jt}) = F_0 \cdot (\bar{\pi}_{jt})$. We parameterize the entrants' amenity distribution as i.i.d log-normal: $\log \delta_j \sim \mathcal{N}(\mu_\delta, \sigma_\delta^2)$ and entry shock as i.i.d. normal: $\eta_j \sim \mathcal{N}(0, \sigma_\eta^2)$.

We now outline the estimation procedure. Recall the entry criterion: firms enter when their value, net of initial debt L_{jt} is larger than equity $sF(\bar{\pi}_{jt})$.

$$\mathbb{E}\left[V(\bar{\pi}_{jt} + \eta_j, (1-s)F_{jt}(1+r_{jt}), \mathbf{S}_{jt})\right] - sF_{jt} > 0$$

with $F_{jt} = F_0 \cdot (\bar{\pi}_{jt} + \eta_j)$ is the initial loan, r_{jt} is the corresponding interest rate.

Estimating F_0 directly through this entry condition requires computing the value function, which needs all the remaining parameters.[35] Instead, we exploit the model's monotonicity to show that this problem is equivalent to estimating a profit threshold for entry.

The threshold equivalence arises from the monotonicity of the value function and our parametric assumptions. For a fixed number of firms n_{jt} (and holding all other parameters fixed), the expression $\mathbb{E}\left[V(\bar{\pi}_{jt} + \eta_j, L_{jt}, \mathbf{S}_{jt})\right] - sF_j$ can be stated as a function of $(\bar{\pi}_{jt} + \eta_j)$.[36] Denote this function as $\mathcal{E}(\pi_j + \eta_j \mid n_{jt}, F_0, \psi_{jt})$. We show that $\mathcal{E}(. \mid n_{jt}, F_0, \psi_{jt})$ is an increasing function: the expected value of the firm, minus its equity, increases with profit. This property implies that, for a given number of initial firms n_{jt}, the above entry criterion is equivalent to a cutoff in profits: firms enter if and only if $\pi_j + \eta_j \geq \pi^*(n_{jt}, F_0, \psi_{jt})$, where $\pi^*(.)$ defines threshold profit levels. (see Appendix 9.10.1).

To improve statistical power, we assume that markets with entry are sufficiently sim-

[35] The Value function depends on the underlying $\tilde{\pi}$, which means it depends on all components of profit (a, D_ϕ, D_ξ, D_C).

[36] Addtionally, the ex-ante type distribution is parameterized to have variance that scales with the mean profit, as shown in the next section.

ilar that variation in n_{jt} has negligible effect on firm value.[37] With this assumption, we pool across all markets with entry, allowing us to estimate a single threshold for entry, which we now call π^*. While strong, this assumption yields a tractable baseline; future extensions will relax it to allow richer market-structure effects.

Now, we specify a likelihood for observed entry $\mathcal{L}(\pi^*, \sigma_\eta, \mu_\delta, \sigma_\delta)$ based on the threshold and the rest of the parameters. The likelihood has two components: contributions from observed entrants and contributions from markets (location-year pairs) with no entry.

Entrants' contribution to the likelihood is simply the probability that they cross the entry threshold: $\Pr(\bar{\pi}_{jt}(\delta_j) + \eta_j \geq \pi^*) = 1 - F_\eta(\pi^* - \bar{\pi}_{jt}(\delta_j))$. For markets with no entry, the contribution is the probability that all potential entrants fall below the threshold, integrating over amenity draws:

$$\Pr(\text{no entry in } m) = \left[\int \underbrace{F_\eta(\pi^* - \bar{\pi}(\delta))}_{\substack{\eta \text{ smaller than} \\ \text{needed for entry}}} dF_{\text{entrant}}(\mu_\delta, \sigma_\delta) \right]$$

We simplify by assuming that incumbent market structure is identical across no-entry markets, using the median market (by number of incumbents) as the reference case.[38]

The full likelihood is:

$$\mathcal{L}(\pi^*, \sigma_\eta, \mu_\delta, \sigma_\delta) = \prod_{\text{entrants } j} \left[1 - F_\eta(\pi^* - \bar{\pi}_{jt}(\delta_j))\right] \times \prod_{\text{no-entry markets } m} \left[\int F_\eta(\pi^* - \bar{\pi}(\delta)) dF_{\text{entrant}}(\mu_\delta, \sigma_\delta) \right] \tag{13}$$

Maximizing this likelihood yields estimates of π^*, σ_η, and $\mu_\delta, \sigma_\delta$.[39]

Identification Identification of π^* comes from balancing the profits of observed entrants and the absence of entry in other markets. If π^* were too low, we would observe entry even in the least attractive markets. If it were too high, we could not explain the

[37]Note that, this only assumes that changes in future firm value due to *future* changes in market structure are similar across market sizes. For this to hold, changes in firm profits from entry or exit should not be too different across markets where entry takes place. We observe this to be a reasonable assumption in our data.

[38]Market characteristics are suppressed to preserve confidentiality.

[39]As econometricians, we do not observe firms' underlying mean profits precisely. Given a prior on ϕ, we could instead integrate over the posterior distribution of firm profits. Our current implementation does not do so, which overstates precision for firms with short panels. Future iterations will incorporate this correction.

presence of low-profit entrants (see Figure 5).

Variance in entry shocks σ_η are identified by the smoothness of the entry margin. Without these shocks, the model would predict a sharp cutoff in entry at exactly the profit level where expected returns equal fixed costs. The observed gradual decline in entry rates as market attractiveness falls identifies the variance of entry shocks needed to smooth this discontinuity.

Finally, the observed distribution of profits among entrants identifies the scale and shape parameter of the entrant-amenity distribution $\mu_\delta, \sigma_\delta$.

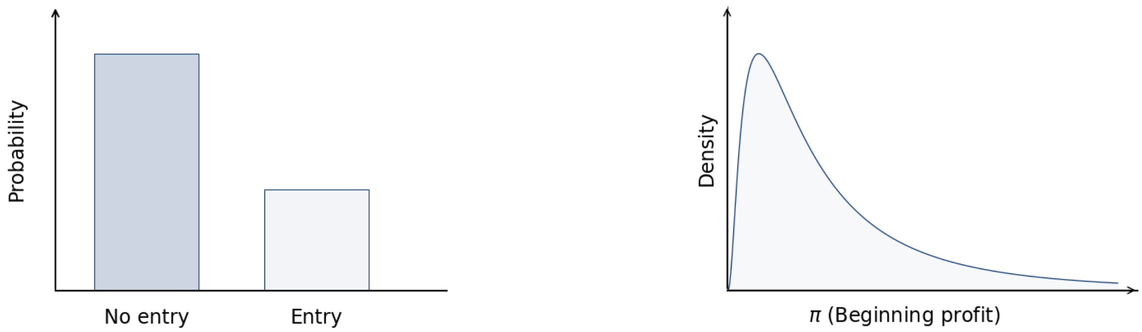

(a) Entry probability across markets

(b) Entrant distribution

Figure 5: Identification from entrants' profits.

Note: Cutoff π^* balances the likelihood of non-exit in city by year pairs (in the left panel) with the likelihood of entry by observed entrants (on the right). η_j explains smoothness in the distribution of entrants. Lognormal parameters fit the observed profit distribution of entrants.

Results: The estimates imply that firms earning roughly \$40,000 per year of net residual profits are threshold entrants. There is wide dispersion in profits among actual entrants, with σ_η estimated at nearly $\$31,050$.

6.2.3 Cost Primitives

We now estimate the remaining primitives of hotel profit streams—(a, D_ϕ, D_C)—along with the entry-cost parameter F_0. These parameters are estimated by maximum likelihood.

The key idea is that a candidate set of parameters for (a, D_ϕ, D_C) dictates the volatility (and mean) of profit streams. Through the lending model, these profit streams translate into firms' exit probabilities via the latent evolution of debt. The estimated parameters, then, are the ones that best fit the observed exit distribution. In parallel, we recover the fixed-cost parameter F_0 using the firms' value function and a correspondence between the threshold π^* and F_0.

Estimation proceeds in three steps. First, we solve the lenders' model offline over a grid of candidate parameters and profits (Section 6.2.3). Second, we use the model to generate exit probabilities for each candidate set of parameters (Section 6.2.3). In practice, we combine these two steps for computational efficiency. Finally, we use these exit probabilities to specify an exit likelihood (Section 6.2.3); within this step, we also estimate F_0 using value functions from the lenders' model.

To implement maximum likelihood, we impose parametric forms on (D_ϕ, D_C). We assume that both are normally distributed: $\pi^{\mathcal{C}} \sim \mathcal{N}(0, \sigma_c^2)$ and $\phi \sim \mathcal{N}(0, \sigma_\phi^2)$.

Solving Lenders' Model We solve the lenders' model to generate the exit probabilities used in constructing the exit likelihood. We solve this offline, over a grid of candidate parameters: (σ_C, σ_ϕ)

To solve the lenders' problem, we need to fully specify the profit stream. Recall that we've already estimated demand deviations π^{ξ}_{jt} in the accommodation model. Choosing a candidate distribution over costs and types then completes the description of profit uncertainty.[40]

The solution of the model yields three sets of objects: interest rates, credit limits and value functions. For a fixed set of candidate parameters and a fixed mean profit, we compute debt limits and interest rates by backward induction from a terminal period. In the final period T, type uncertainty is fully resolved and the model is stationary. Using debt limits in T, we compute debt limits and interest rates in $T-1$. We continue through backwards induction from $T-1$ to period 0, calculating debt limits and interest rates at each step. At each step, we use the lenders' age-specific posterior distribution over firm types (see Appendix 9.10.1). For estimation, we set $T = 10$ and verify that results are robust to increasing the horizon.

We solve the model over a full grid of candidate parameters and profits. Firm profits are discretized over their observed support. For computational tractability, we assume that there are three types whose dispersion is obtained from a three node Gaussian quadrature. Similarly, we discretize over a wide range of candidate cost and type dispersions (σ_C, σ_ϕ). Implementation details in the Appendix 9.10.3.

With interest rates and credit limits in hand, we solve for each firms' value functions. The value function satisfies the Bellman equation (6), which we solve as a linear system rather than by direct iteration. This approach—similar to Hotz and Miller (1993a)—exploits the linear structure of the expected-value operator and reduces run-

[40]The mortgage parameters simply rescales mean profit, so we don't need to solve the model separately for different values of a.

time by roughly two orders of magnitude (see Appendix 9.10.3). Finally, we account for market-structure dynamics by averaging value functions over transitions in the competitive environment $\psi_{jt}(.)$. Further details are provided in the Appendix 9.10.3.

Obtaining Exit Probabilities In this step, we combine interest rates, credit limits, and candidate profit uncertainty to compute firm-level exit probabilities.

The key idea is that interest rates and realized profits together determine how a firms' debt evolves. Recall the debt rollover equation (4) from the lenders' model:

$$L_{jt+1} = \max\{L_{jt}(1 + r_{jt}) - \pi_{jt}, 0\}$$

For any t liability L_{jt} , we can derive the probability of reaching debt L_{jt+1}, based on the distribution of the profit stream π_{jt}. We denote this one-period mapping as the transition kernel $K_{jt}(L_{jt+1} \mid L_{jt})$. In our discrete implementation K_{jt} is simply a transition matrix over debt states.

By iterating this kernel forward, we construct a multi-period—or "long-run"—transition matrix, $\Gamma_j(L_{t+k} \mid L_t)$, which gives the distribution of liabilities k periods ahead in time. A key computational advantage is that these transition matrices can be generated recursively while solving the lenders' model (see Appendix 9.10.3 for further details).

Starting from any liability in period t, the probability that a firm exits in period $t + k$ is given by this long run transition kernel: specifically, the probability that debt first exceeds the equilibrium credit limit in period $t + k$ conditional on survival until $t + k - 1$.

Let $p_{t+k}(L_t|\bar{\pi}_j)$ denote the probability that firm j with mean profit $\bar{\pi}_j$ exits in period $t + k$, starting from debt L_t in period t[41]. Let $q_{t+k}(L_t \mid \bar{\pi}_j)$ denote the probability that the firm survives until the end of period $t + k$. These two objects—exit and survival probabilities—form the foundation of the likelihood function estimated in the next step.

Maximizing Exit Likelihood We use the model-implied exit probabilities to construct a likelihood for exit based on a set of candidate parameters $(a, \sigma_\phi, \sigma_C)$.

The procedure has two stages.First, we determine firms' starting debt levels, since exit probabilities depend on initial debt positions. We obtain these starting debts by estimating the fixed-cost parameter F_0. Second, we use the resulting debt levels to compute the likelihood of exit and survival.

[41]Note here we're denoting firm profits as fixed, as assumed in 6.1.2. This is not necessary, the transition matrices provide an elegant way to translate debt across periods even when mean profits shift. In practice, this means that instead of using exit probabilities indexed by profit, we need to use exit probabilities indexed by profit history. This increases runtime.

We recover F_0 using the estimated profit threshold π^*. For each candidate (σ_ϕ, σ_C), , we have already computed the value function while solving the model. The optimal F_0 is that one that best rationalizes π^* as the entry threshold, given the other entry parameters $(\eta, \mu_\delta, \sigma_\delta)$. We simplify this computation by noting that the conditional likelihood is maximized exactly when the threshold firm is indifferent to entry (see Appendix 9.10.1). The likelihood maximand, therefore, is simply the F_0 that solves the indifference condition:

$$\mathbb{E}\left[V(\pi^*, (1-s)\pi^* F_0(1+r_*), \mathbf{S}_{jt})\right] - s\pi^* F_0 = 0$$

where r_* is the corresponding interest rate on starting debt. We calibrate the borrower equity share to $s = 0.1$, consistent with SBA requirements of 5–20% equity contribution in guaranteed loans (*SBA SOP* 2023, *Code of Federal Regulations* 2024), since there is limited variation to identify this parameter in our data.

After estimating F_0 the initial debt for entrants is given by $L_{jt} = (1-s)F_0 \cdot \bar{\pi}_{jt}$. To define a likelihood, we need an initial debt for incumbents as well. For these firms, we assign the distribution of debt observed $t+k$ periods after entry as this starting debt, where k is the mean age of incumbents[42].

We then use exit and survival probabilities conditional on starting debt to form the likelihood. Let τ_j denote the time between the first and the last appearance of j in the panel. Let E_j be an indicator for firm exit, $\bar{\pi}_{jt}$ be mean profit and L_{jt} denote the firm's initial debt. The likelihood takes the following expression:

$$\mathcal{L}(\sigma_\phi, \sigma_c, a) = \Pi_j \left[p_{\tau_j}(L_{jt} \mid \bar{\pi}_{jt}) \right]^{E_j} \left[q_{\tau_j}(L_{jt} \mid \bar{\pi}_{jt}) \right]^{1-E_j}$$

Here, $p_{\tau_j}(L_{jt} \mid \bar{\pi}_{jt})$ is the probability of exit for firm j after τ_j periods, and $q_{\tau_j}(L_{jt} \mid \bar{\pi}_{jt})$ is the probability of surviving for τ_j periods without exit, both conditional on starting debt L_j and mean profit $\bar{\pi}_j$. We maximize this likelihood with respect to $(\sigma_\phi, \sigma_c, a)$ (see Appendix 9.10.3 for details).[43]

Identification The profit-stream parameters a, D_ϕ and D_C are identified from the joint distribution of firm profits and exit. Mortgage payments a and the distribution of costs D_C are identified from the cross-sectional pattern of exit. The distribution of hidden types D_ϕ is identified from panel variation in exit.

[42] k suppressed to preserve disclosure protection

[43] As econometricians, we do not observe firms' mean profits exactly. For a given σ_ϕ, we can derive a posterior distribution over each firm's mean profit. In a prior specification without mortgages, we compared versions that did and did not account for this posterior; the resulting point estimates were similar. Future work will extend this correction to the full model.

We begin with mortgage share a. Recall that mortgage payments reduce a firm's mean profit from $\bar{\pi}_{jt}$ to $\bar{\pi}_{jt}(1-a)$ without affecting uncertainty $\tilde{\pi}_{jt}$. Lower levels of profit relative to the uncertainty increase firm volatility. As a result, higher a makes all firms more volatile, raising overall exit probability. We identify a from average exit rates across the sample.

Cost shocks $\pi^{\mathcal{C}}_{jt}$ rationalize higher exit among less profitable firms. Since σ_C is common across all firms, it affects these firms more relative to firms with higher profit. Figure (6) shows this empirical pattern: firms with lower profits exit more frequently. The estimated variance of costs is the one that best matches exit probability across profits with its empirical distribution[44].

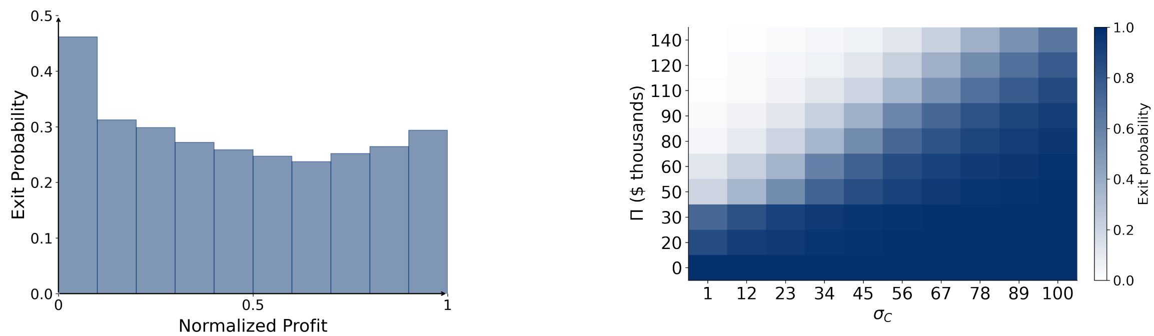

Figure 6: Empirical vs. model-implied firm exit probabilities.

Note: The left plot shows the empirical distribution of firm exit, with profit on the x axis and exit probability on the y-axis. The right figure shows a heatmap of model implied exit probabilities, with $\sigma_{\tilde{\pi}}$, variance on profit innovations on the x-axis, profit on the y-axis and darker colors representing more exit. As $\sigma_{\tilde{\pi}}$ increases, exit probability increases secularly, but small firms always exit more. The estimated value of σ_c and a matches the empirical exit probability on the left to the model implied one. Estimates cleared for disclosure by U.S. Census under disclosure ID 12761.

Hidden quality types D_ϕ are identified from age-related patterns in exit. Young firms face higher exit risk because lenders haven't yet learned their true quality, leading to more dispersion and tighter credit constraints. As firms age and reveal information through profit histories, their exit risk declines. Figure (7b) illustrates this mechanism: when type uncertainty is high ($\sigma_\phi > 0$), the model predicts larger differences in exit rates between young and old firms. The estimated ϕ matches the empirical exit gradient to the one implied by the model.

Results: The estimates imply substantial uncertainty in hotel markets. Hidden type uncertainty is economically significant—the estimated standard deviation equals 60% of mean profits—consistent with empirical evidence on young firm riskiness. Idiosyncratic

[44]Our framework allows for non-parametric identification of $D_{\mathcal{C}}$ under mild regularity conditions (see Appendix 9.11), though we adopt a parametric specification for estimation power.

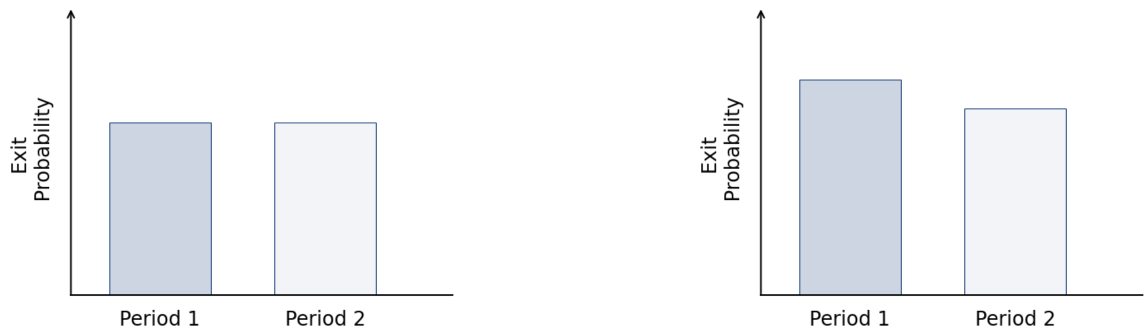

(a) $\sigma_\phi = 0$ (b) $\sigma_\phi > 0$

Figure 7: Age Exit Gradient with different ϕ

Note: Left plot shows exit probability for a firm with given distribution of π' across two periods when $\sigma_\phi = 0$, i.e., there is no type uncertainty. The right plot shows higher exit probability for a firm in its first period compared to the next period, because lenders resolve uncertainty over time. The estimate σ_ϕ matches difference in exit probability implied by the model to its empirical (initial condition adjusted) counterpart.

cost shocks are estimated to have a standard deviation around $3,400 per year, which is about 60% of profit of our smallest 5% of the firms. These shocks disproportionately affect smaller firms, as a given dollar shock represents a larger fraction of their revenue base. The interaction between cost uncertainty and firm size helps explain the observed higher exit rates among smaller establishments.

Estimated proportion of mortgage payments is low and insignificant at 4% of profit, but in conjugation with the estimated marginal costs of more than 75% (which partially captures pricing wedges to drive down leverage)[45], firms pay nearly 80%t of revenue towards their costs and debt, which aligns with industry benchmarks and IRS reports.

Finally, firms are deeply financially constrained due to recovery risk. The estimated debt limits for working capital loans are between three to five times yearly variable profit. Compare, this to the firm's discounted lifetime value. At the chosen discount factor of $\beta = 0.96$ this is 25 times the mean profit.

6.3 Model fit

The model closely matches the observed distribution of firm dynamics, as shown in Figure (8).

As a robustness check, introducing an additional exogenous exit rate yields insignificant estimates and does not improve fit, suggesting that the parsimonious structure is

[45] See Appendix 9.8.5 for a discussion of why firms might price as if their marginal costs are higher than their actual per unit cost in this setting.

Figure 8: Exit Probability in Hotels

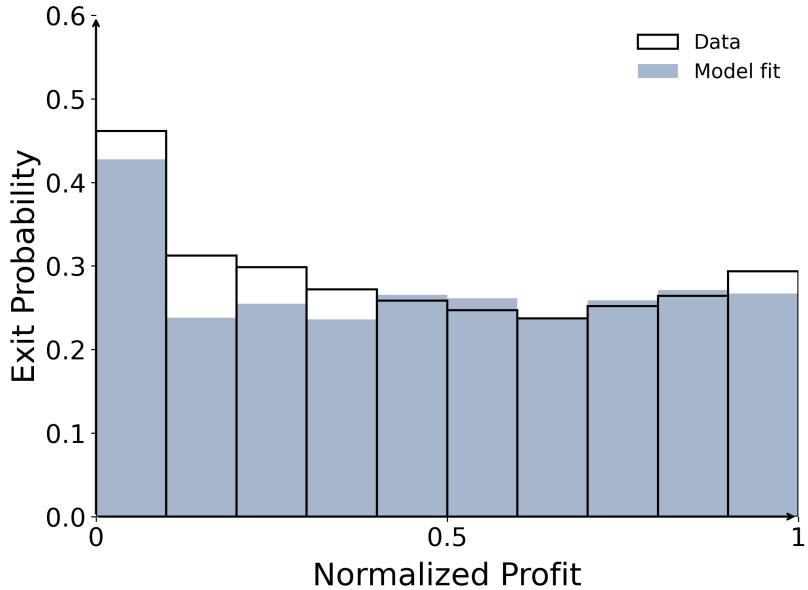

Note: This figure plots the implied exit probability from the model against empirical exit probability in the 10 year period. Estimates cleared for disclosure by U.S. Census under disclosure ID 12761.

sufficient to explain observed exit patterns.

Forward simulations of the model, in which we solve for equilibrium prices, entry, and exit, also generate outcomes consistent with the data: the implied inter-quartile range of prices is 83–97, well within the range observed in the accommodation market.

Together, these checks provide confidence that the model provides a credible basis for counterfactual analysis.

7 Counterfactuals

We use our estimated model to run counterfactual experiments. These enable us to understand the gains from 7(a) awards, and the channels through which they operate.

First, we use our model estimates to simulate markets with and without 7(a) assistance. This allows us to quantify the welfare effects of the 7(a) program. Our counterfactuals imply that hotel markets in our sample are inefficiently constrained by financial frictions, and alleviating these constraints raises welfare. Program awards from 7(a) create these gains cost-effectively. Total welfare increases by about \$54 million (0.6% of market revenue), at a fiscal cost of \$24 million in guarantee payments, implying a fiscal multiplier of 2.25.

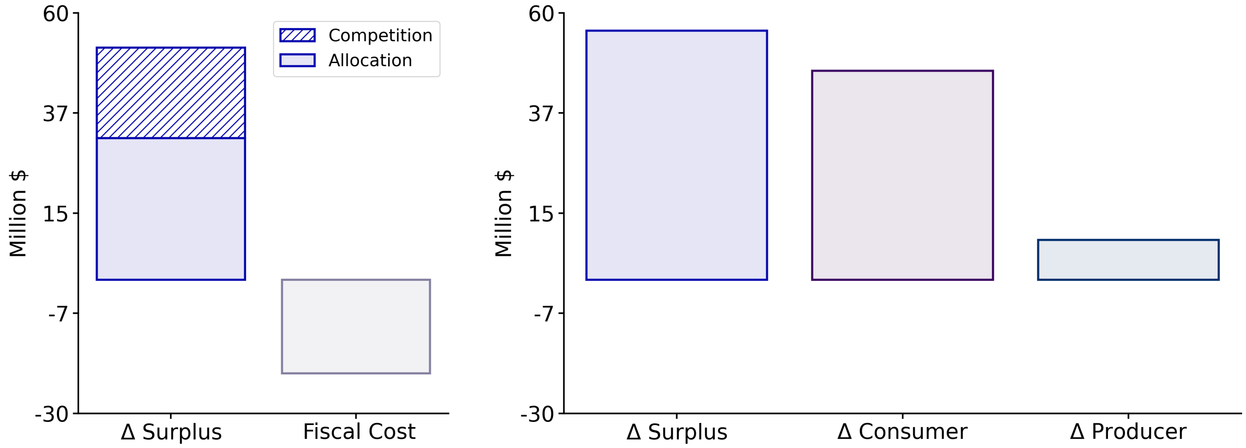

Figure 9: Welfare and Fiscal Cost of 7(a)

Note: The left plot shows the overall surplus, split into the value of additional, efficient firms (allocation) and the surplus created by reduction in prices (competition). Fiscal costs are smaller compared to the overall gains. The right plot breaks down the overall surplus gains into gains of consumers and hotels. Estimates cleared for disclosure by U.S. Census under disclosure ID 12761.

These gains come from increased hotel presence in our markets, along with decreases in hotels' debt expenses. Receiving a 7(a) award substantially reduces a hotel's borrowing costs and increases its credit limit. This credit expansion decreases or delays exit by beneficiary hotels. Cheaper credit and longer survival also makes hotels that receive 7(a) awards more valuable, and this encourages new entry. Markets affected by 7(a) feature 1.4% more firm-years.

We decompose welfare into two pieces—increased allocation to valuable firms and competition. First, we shut down price changes and evaluate welfare effects. This reflects how program-led expansion of hotel options changes welfare in the markets. We find that, on net, these additional firms expand the market, delivering accommodation services that were inefficiently prevented due to credit frictions. A 7(a) award increases credit allocation to these positive value firms. Roughly $31 million of total gains are attributable to this channel.

Competition effects further increase surplus. As more firms enter and survive, competition intensifies, lowering prices. Consumers, who are highly price sensitive, gain substantially from these price reductions. This more than offsets losses to non-beneficiary hotels, whose profits decrease. On average, prices fall by 0.6% across markets, with upto 20% declines in the most concentrated markets. Price effects generate an additional $23 million in welfare gains. Figure (9) illustrates this breakdown. Overall, the quantity of accommodation sold increases by nearly 1% from both of these channels.

Next we understand how these gains are distributed (see Figure 9). Consumers capture most of the surplus, gaining $46 million overall. Nearly half of this comes from lower prices, while the remainder comes from having more options. Consumer gains are highly uneven across markets. Markets where 7(a) induces new entry—representing just 8% of aggregate revenue—generate 25% of total consumer gains. Even within these markets, the distribution is skewed: in the most concentrated decile (by Herfindahl–Hirschman Index), consumer surplus rises by more than 10% of baseline revenue, compared to a program-wide average gain of 0.5%.

Hotels gain $8 million on net, but with considerable heterogeneity. We decompose effects based on whether 7(a) support induces new (marginal) entry or not. In markets where a 7(a) award introduces a new hotel, incumbents lose profits due to significant business stealing, sometimes even causing exit. After accounting for entry costs, entrants themselves earn meager increases in producer surplus. The net producer surplus in these markets is -$0.2 million: from a $1.4 million increase of entrant surplus, and a drop of nearly $1.6 million of incumbent surplus.

Hotels that receive 7(a) awards, but are not marginal entrants, create welfare cost-effectively. A 7(a) award reduces such a hotel's borrowing costs and decreases its exit probability, which increases dividends. These changes do not impose large fiscal costs—by preventing firm exit, the program lowers its guarantee payments. For example, even for the marginal firm that earns $38,000 in profit, a 7(a) guarantee lowers its one-period exit probability from 0.19 to 0.09. This feature minimizes program costs while increasing hotel dividends through cheaper credit.

7.1 Alternate Targeting

To study how program design might be improved, we examine reallocation and optimal targeting. First, we conduct a simple exercise where we shift awards to concentrated markets. Specifically, we remove all awards from the least concentrated third of the markets (by HHI), and double the award probability for the most concentrated third of the markets. Relative to a no-SBA benchmark, consumers gain $52 million, hotels gain $5 million, and the government pays $30 million in guarantees, resulting in a lower multiplier (relative to current 7(a) awards) of nearly two. This aligns with our baseline finding: awards to concentrated markets expand consumer welfare through entry, while producer gains are limited due to substantial business-stealing effects.

Next, we try to understand how these awards can be optimally targeted. Specifically, we ask: if the government decides an (ex-ante) award probability based on candidate hotels' amenity value and the corresponding market's composition, what would the optimal scheme resemble? While we formally describe this problem, and theoretically illustrate the "first-order" like conditions of such an optimum, we're unable to make empirical progress due to significant computational challenges. (see Appendix 9.11)

Instead, we focus on a simpler allocation problem: government chooses award probability based on a hotels' profit (proxy for its characteristics), and the HHI of the target market (proxy for market structure). For this exercise, we divide markets into 3 percentile bins by HHI, and firms into 5 percentile bins by profit. We restrict attention to binary allocation schemes, where social planner gives awards with probability one or zero. This lends maximum power to our counterfactuals to illustrate program effectiveness, measured through the fiscal multiplier. We think of this exercise as conceptually illustrative, rather than a policy prescription. We focus on two versions of this problem: one where the government places equal weights on producers and consumers, and one where it only values consumer surplus.

Optimal allocations in either scenarios avoid hotels in the smallest 20th profit bin. These firms are typically young entrants, and default with high probability: operational costs are highly damaging to small hotels in their early, indebted years. While they can create significant value in concentrated markets, the government often needs to pay guarantees when they exit, in addition to the hotels' private entry cost. Any gains created over their short life-span are not cost effective, and excluding them significantly increases the fiscal multiplier.

Awards don't substantially impact the largest 80-100th percentile firms on the extensive margin. Most of these firms are incumbents with low debt. Exit probability is low, especially after receiving 7(a) awards, which makes fiscal costs calculation error-prone [46]. For now, we conclude that these firms are not marginal to the policy and exclude them from this exercise, with the caveat that we lack accurate measures of program multiplier.

Both the efficient and consumer-optimal allocations award firms in the 20-80 percentile of profits, but in separate markets. Efficient allocations target the least concentrated markets. These markets have many firms, and as such, 7(a) rarely induces new entry. This is because business stealing effects, along with the large number of firms, dictate that a new hotel needs to have substantial amenity value to enter. Typically, entry (and marginal entry) happens only in the wake of exit. Therefore, most 7(a) awards go to existing firms. As noted before, such awards create additional dividends for hotels

[46] Exit occurred in less than 10% of our simulations.

through reduced borrowing costs and decreased exit. Additionally, changes in market structure due to churn contribute little to the firms' profit volatility (due to the large number of firms in these concentrated markets). Our simulations find similar multipliers within the 20-80 percentile bin, so we quote results from a counterfactual that jointly awards these bins with probability one to minimize simulation error.[47] Under this allocation, hotels gain \$130 million, consumers gain \$45 million and the government pays \$40 million as guarantees, leading to a multiplier of 4.38.

On the contrary, firms enter concentrated markets more often, and 7(a) induces marginal entry. When an additional hotel enters, consumers benefit from price drops and increase in the set of hotel options. There are some gains in producer welfare— with incumbents losing profits on net and entrants gaining it, but all hotels benefit from reduction in borrowing costs. Entry, however, is still privately costly. Additionally, new entry substantially increases the chances of incumbent exit relative to markets with larger firms. Finally, these markets usually have lower market sizes (Bresnahan and Reiss 1991), and firm profits as such tend to be smaller. This means that cost shocks are more impactful, which additionally increases churn in these markets, which is costly. Under this allocation consumers gain \$150 million, hotels gain \$35 million and the government pays \$45 million in guarantees, leading to a multiplier of 4.1. Figure (fig:fourpanels) visualizes the allocations under various counterfactual schemes.

Overall, these counterfactuals highlight that welfare effects of 7(a) allocations depend crucially on program targeting. While the multipliers in each of these cases are comparable, there are large differences in who gains from either of these allocations.[48]

7.2 Discussion

Our counterfactual exercises underscore broader insights about credit policy in markets with imperfect competition.

First, credit frictions shape market structure and has important welfare consequences. In our setting, small firms rely solely on debt, and debt is risky from imperfect recovery.

[47] This approach ignores spillovers across bins. We notice slight changes in multipliers from this joint exercise compared to individual averages, but fail to reject them based on dispersion. At our current number of simulations, we consider it reasonable to club these bins, treating them as a unit and reducing simulation error through more award observations.

[48] A general caveat on multipliers from these high probability allocations is that firm exit events may lie beyond the simulation length. While we always account for this by using a discounted mean cost in the last period summed over periods 11 to 20, with very large award probabilities, 10 years maybe insufficient to accurately account for firm exit for later periods. In future iterations, we shall attempt to capture this using longer runs.

Figure 10: Award Schemes across Counterfactuals

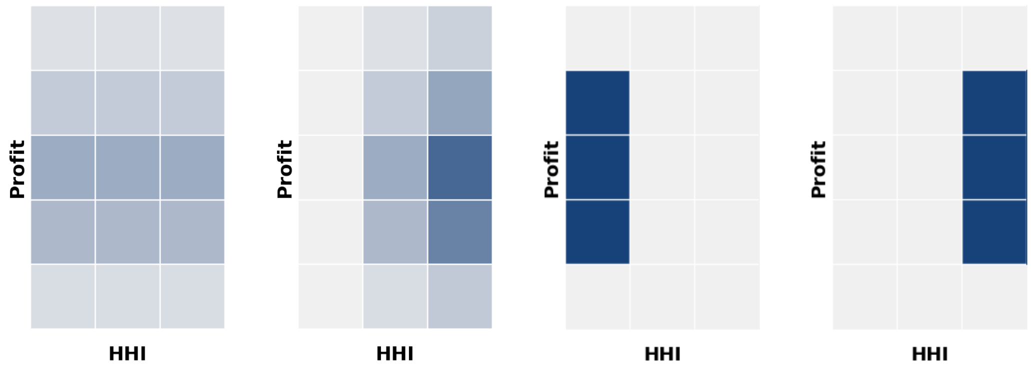

Note: This figure shows the heatmap of 7(a) awards in various counterfactuals, where HHI terciles are on the x-axis and profit quintiles are on the y-axis. Panel (1) from the left shows the baseline 7(a) probability. Panel (2) shows our simple reallocation where we remove allocations from least concentrated HHI tercile and allocate at a double rate within the most concentrated HHI tercile. Panel (3) represents the equi-weighted optimal allocation, which awards 20-80 profit percentile firms with probability 1.0. Panel (4) represents the consumer optimal allocation, which awards the 20-80 percentiles with probability 1.0.

The combination of these restricts firms' access to financing and meaningfully contracts markets.

Second, credit guarantees, like the SBA 7(a), provide a cost effective way to expand markets and increase welfare. In our setting, the insurance offered by 7(a) substantially increases access to credit for small firms. Payments due to default are limited, as awards themselves reduce exits, and the government is only partially liable. This results in large multiplier.

Third, targeting matters significantly for both efficiency and distribution. In our setting, we find program support to the smallest of small firms to be substantially less cost effective. Depending on whether the planner targets more or less concentrated markets, consumers or hotels are the main beneficiaries.

These insights extend beyond the 7(a) program to a wider set of credit policies that influence firm financing and market structure. Our framework suggests that credit interventions are most effective in sectors with substantial recovery issues, significant ex-ante uncertainty, binding credit frictions, and concentrated product markets—characteristics common in small-business-intensive industries like restaurants and local medical and dental practices.

8 Conclusion

In this paper, we develop and estimate a structural model linking debt and accommodation markets. This allows us to quantify the welfare effects of SBA's 7(a) loan guarantee program. Our approach endogenizes both lender behavior and market competition, which enables us to run credible credible counterfactual experiments.

We find that 7(a) generates substantial welfare gains of $56 million at a fiscal cost of $24 million, yielding a multiplier greater than two. These gains operate through two channels: allocative efficiency improvements ($31 million) that direct new credit toward high-value firms, and increased competition effects ($23 million) that lower prices and expand consumer choice. Consumers capture the majority of benefits despite not being direct program participants, highlighting the importance of equilibrium effects in policy evaluation.

Our counterfactual exercises demonstrate that program design significantly affects both the magnitude and distribution of gains. Optimal targeting schemes can benefit consumers or producers, depending on which markets they focus on.

Our framework provides a tractable approach to modeling financial policy and firm dynamics, though it relies on several simplifying assumptions. We treat default and exit as equivalent events and assume lender market structure remains fixed. The first assumption may be less appropriate in industries where exit stems from strategic or operational decisions unrelated to financial distress. The second becomes less compelling when policies are large enough to reshape lender-side industry structure.

Future research could extend this analysis by endogenizing lender market structure to understand how large-scale programs affect credit supply. Additionally, applying the framework to settings with clean default and exit data could help assess the content of our assumptions, while testing the approach across different small business sectors would provide valuable evidence on external validity. These extensions would enhance our understanding of credit policy effects across diverse economic environments.

References

Adams, William, Liran Einav, and Jonathan Levin (Mar. 2009). "Liquidity Constraints and Imperfect Information in Subprime Lending". In: *American Economic Review* 99.1, pp. 49–84.

American Hotel & Lodging Association (Apr. 2020). *Letter to Congressional Leadership on CARES Act and SBA Loan Limits*. Requests increase in SBA loan funding and update to CARES Act to help hoteliers manage payroll and debt.

— (2025). *2025 State of the Hotel Industry Report*. Tech. rep. Accessed September 30, 2025. AHLA.

Bachas, Natalie, Olivia S. Kim, and Constantine Yannelis (2021). “Loan guarantees and credit supply”. In: *Journal of Financial Economics* 139.3, pp. 872–894.

Banerjee, Abhijit V. and Esther Duflo (None 2014). “Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program”. In: *The Review of Economic Studies* 81.2, pp. 572–607.

Bazzi, Samuel et al. (Sept. 2023). *Credit Supply Shocks and Firm Dynamics: Evidence from Brazil*. Working Paper 31721. National Bureau of Economic Research.

Benmelech, Efraim and Nittai K. Bergman (2009). “Collateral Pricing”. In: *Journal of Financial Economics* 91.3, pp. 339–360.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). “Automobile Prices in Market Equilibrium”. In: *Econometrica* 63.4, pp. 841–890.

Berry, Steven and Joel Waldfogel (1999). “Advertising, Geography, and Aggregation: The Case of Cable Television”. In: *Economica* 66.262, pp. 35–50.

Berry, Steven T. (1994). “Estimating Discrete-Choice Models of Product Differentiation”. In: *The RAND Journal of Economics* 25.2, pp. 242–262.

Bresnahan, Timothy F and Peter C Reiss (1991). “Entry and competition in concentrated markets”. In: *Journal of political economy* 99.5, pp. 977–1009.

Bresnahan, Timothy F., Scott Stern, and Manuel Trajtenberg (1997). “Market Segmentation and the Sources of Rents from Innovation: Personal Computers in the Late 1980s”. In: *RAND Journal of Economics* 28, S17–S44.

Brown, J. David and John S. Earle (2017). “Finance and Growth at the Firm Level: Evidence from SBA Loans”. In: *The Journal of Finance* 72.3, pp. 1039–1080. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12492`.

Brown, J. David et al. (2024). *Financing, Ownership, and Performance: A Novel, Longitudinal Firm-Level Database*. Working Paper 33254. National Bureau of Economic Research.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin (Aug. 2011). “Finance and Development: A Tale of Two Sectors”. In: *American Economic Review* 101.5, pp. 1964–2002.

Buera, Francisco J., Benjamin Moll, and Andrei Shleifer (2013). “The Macroeconomics of Microfinance”. In: *NBER Working Paper No. 17905*. Also published in Journal of Economic Perspectives (2015), 29(1): 85–112.

Bulow, Jeremy and Kenneth Rogoff (1989). “A Constant Recontracting Model of Sovereign Debt”. In: *Journal of Political Economy* 97.1, pp. 155–178.

Cardell, N. Scott (1997). “Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity”. In: *Econometric Theory* 13.2, pp. 185–213.

Cavalcanti, Tiago V et al. (Apr. 2021). *Dispersion in Financing Costs and Development*. Working Paper 28635. National Bureau of Economic Research.

Chevalier, Judith A. (June 1995). “Capital Structure and Product-Market Competition: Empirical Evidence from the Supermarket Industry”. In: *American Economic Review* 85.3. Stable URL included, pp. 415–435.

Chevalier, Judith A. and David S. Scharfstein (Sept. 1996). “Capital-Market Imperfections and Countercyclical Markups: Theory and Evidence”. In: *American Economic Review* 86.4. Also circulated as NBER Working Paper No. 4614 (1994)., pp. 703–725.

Chodorow-Reich, Gabriel (2014). “The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008–9 Financial Crisis”. In: *Quarterly Journal of Economics* 129.1, pp. 1–59.

Code of Federal Regulations (2024). Code of Federal Regulations. 3 CFR §120.910 — What is the Borrower’s Contribution?.

Crawford, Gregory S., Nicola Pavanini, and Fabiano Schivardi (July 2018). “Asymmetric Information and Imperfect Competition in Lending Markets”. In: *American Economic Review* 108.7, pp. 1659–1701.

Cushman & Wakefield (May 2018). *Focus on Hotel: Construction Costs — Global Hospitality*. Tech. rep. Accessed September 30, 2025. Cushman & Wakefield.

Davis, Steven J. and John C. Haltiwanger (1992). “Gross Job Creation, Gross Job Destruction, and Employment Reallocation”. In: *The Quarterly Journal of Economics* 107.3, pp. 819–863.

Dilger, Robert Jay (2019). *Small Business Administration and Job Creation*. CRS Report R41523. Congressional Research Service.

Dilger, Robert Jay and Sean Lowry (Mar. 2020). *Small Business Administration 7(a) Loan Guaranty Program*. Tech. rep. R41146. Congressional Research Service.

Dixit, Avinash K. and Albert S. Kyle (1985). “The Use of Protection and Subsidies for Entry Promotion and Deterrence”. In: *American Economic Review* 75.1, pp. 139–152.

Doraszelski, Ulrich, João F. Gomes, and Felix Nockher (2022). “Oligopoly Dynamics With Financial Frictions”. SSRN Working Paper, posted November 19, 2022.

Dou, Winston Wei, Wei Wang, and Wenyu Wang (2022). "The Cost of Intermediary Market Power for Distressed Borrowers". In: *Working paper / forthcoming*. Available at SSRN.

Dunne, Timothy et al. (2013). "Entry, Exit, and the Determinants of Market Structure". In: *RAND Journal of Economics* 44.3, pp. 462–487.

Eaton, Jonathan and Mark Gersovitz (1981). "Debt with Potential Repudiation: Theoretical and Empirical Analysis". In: *The Review of Economic Studies* 48.2, pp. 289–309.

Einav, Liran, Michael Jenkins, and Jonathan Levin (2012). "Title Unknown — need full details". In: *Unknown*. Please provide full reference; I could not definitively identify this one.

Eisfeldt, Andrea L. and Adriano A. Rampini (2009). "Leasing, Ability to Repossess, and Debt Capacity". In: *Review of Financial Studies* 22.4, pp. 1621–1657.

Fan, Ying, Kai-Uwe Kühn, and Francine Lafontaine (2017). "Financial Constraints and Moral Hazard: The Case of Franchising". In: *Journal of Political Economy* 125.6, pp. 2082–2125. eprint: `https://doi.org/10.1086/694566`.

Fan, Ying and Mo Xiao (2015). "Competition and Subsidies in the Deregulated US Local Telephone Industry". In: *RAND Journal of Economics* 46.4, pp. 751–776.

Farronato, Chiara and Andrey Fradkin (June 2022). "The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry". In: *American Economic Review* 112.6, pp. 1782–1817.

Federal Reserve Banks (Apr. 2019). *Small Business Credit Survey: 2019 Report on Employer Firms*. Based on 2018 survey data; covers financing needs, performance, including reliance on credit for operating expenses.

Gandhi, Amit and Jean-François Houde (2019). *Measuring Substitution Patterns in Differentiated-Products Industries*. Working Paper 26375. Revised version available. National Bureau of Economic Research.

Geanakoplos, John (1997). "Promises, Promises". In: *The Economy as an Evolving Complex System II*. Ed. by W. Brian Arthur, Steven N. Durlauf, and David A. Lane. Westview Press, pp. 285–320.

Geelen, Thomas (2016). "Debt Maturity and Lumpy Debt". In: *SSRN (preprint)*. Available at SSRN.

Geelen, Tim A. (2024). "Asset Life, Leverage, and Debt Maturity Matching". In: *The Journal of Finance*. Forthcoming.

Graham, John R. and Campbell R. Harvey (2001). "The Theory and Practice of Corporate Finance: Evidence from the Field". In: *Journal of Financial Economics* 60.2–3, pp. 187–243.

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen (Feb. 2020). “Do Credit Market Shocks Affect the Real Economy? Quasi-experimental Evidence from the Great Recession and "Normal" Economic Times”. In: *American Economic Journal: Economic Policy* 12.1, pp. 200–225.

Gu, Zheng and Luyuan Gao (2000). “A Multivariate Model for Predicting Business Failures of Hospitality Firms”. In: *Tourism and Hospitality Research* 2.1, pp. 37–49. eprint: https://doi.org/10.1177/146735840000200108.

Hart, Oliver and John Moore (1994). “A Theory of Debt Based on the Inalienability of Human Capital”. In: *The Quarterly Journal of Economics* 109.4, pp. 841–879.

Heckman, James J., Hidehiko Ichimura, and Petra Todd (1998). “Matching as an Econometric Evaluation Estimator”. In: *The Review of Economic Studies* 65.2, pp. 261–294.

Hennessy, Christopher A. and Toni M. Whited (2005). “Debt Dynamics”. In: *The Journal of Finance* 60.3, pp. 1129–1165.

— (2007). “How Costly Is External Financing? Evidence from a Structural Estimation”. In: *The Journal of Finance* 62.4, pp. 1705–1745. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2007.01255.x.

Hollenbeck, Brett (2017). “The Economic Advantages of Chain Organization”. In: *RAND Journal of Economics* 48.4. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2621840, –.

Holmström, Bengt and Jean Tirole (1997). “Financial Intermediation, Loanable Funds, and the Real Sector”. In: *The Quarterly Journal of Economics* 112.3, pp. 663–691.

Hortaçsu, Ali et al. (2019). “Indirect Costs of Financial Distress: Fire Sales in Auto Lending”. In: *American Economic Review* 109.7, pp. 2463–2508.

Hospitality Financial and Technology Professionals (2024). *Navigating Hotel Financing in a World Eager for New Travel Experiences*. Tech. rep. HFTP.

Hotz, V. Joseph and Robert A. Miller (1993a). “Conditional Choice Probabilities and the Estimation of Dynamic Models”. In: *The Review of Economic Studies* 60.3, pp. 497–529.

— (1993b). “Conditional Choice Probabilities and the Estimation of Dynamic Models”. In: *The Review of Economic Studies* 60.3, pp. 497–529.

Jermann, Urban and Vincenzo Quadrini (2012). “Macroeconomic Effects of Financial Shocks”. In: *American Economic Review* 102.1, pp. 238–271.

Jovanovic, Boyan (1982). “Selection and the Evolution of Industry”. In: *Econometrica* 50.3, pp. 649–670.

Kastl, Jakub (2012). *Empirical Finance and Banking (Working Paper)*. Tech. rep. Working paper – see “Empirical Finance” folder on author’s site. Princeton University / personal website.

Kehoe, Timothy J. and David K. Levine (1993). “Debt-Constrained Asset Markets”. In: *Review of Economic Studies* 60.4, pp. 865–888.

Kermani, Amir and Yueran Ma (2020). *Two Tales of Debt*. NBER Working Paper 27641. Revised version cited in text. National Bureau of Economic Research.

Kiyotaki, Nobuhiro and John Moore (1997). “Credit Cycles”. In: *Journal of Political Economy* 105.2, pp. 211–248.

Kocherlakota, Narayana R. (1996). “Implications of Efficient Risk Sharing without Commitment”. In: *Review of Economic Studies* 63.4, pp. 595–609.

Lelarge, Claire, David Sraer, and David Thesmar (Feb. 2010). *Entrepreneurship and Credit Constraints: Evidence from a French Loan Guarantee Program*.

Lodging Staff (Mar. 2022). *Lodging's 2022 Guide to Lenders and Financial Service Providers*. Lodging Magazine (Official Publication of AHLA). Accessed September 30, 2025.

Ma, Yueran and José A. Scheinkman (2020). *Going-Concern Debt of Financial Intermediaries*. NBER Working Paper w28088. Revised 2023. National Bureau of Economic Research.

Mankiw, N. Gregory and Michael D. Whinston (1986). “Free Entry and Social Inefficiency”. In: *RAND Journal of Economics* 17.1, pp. 48–58.

Metaxas, Theodore and Athanasios Romanopoulos (2023). “A Literature Review on the Financial Determinants of Hotel Default”. In: *Journal of Risk and Financial Management* 16.7.

Midrigan, Virgiliu and Daniel Yi Xu (Feb. 2014). “Finance and Misallocation: Evidence from Plant-Level Data”. In: *American Economic Review* 104.2, pp. 422–58.

Miller, Nathan H. and Matthew C. Weinberg (2017). “Understanding the Price Effects of the MillerCoors Joint Venture”. In: *Econometrica* 85.6, pp. 1763–1791.

Mills, Karen G., Brayden McCarthy, and Christopher F. Mayer (2014). “The State of Small Business Lending: Credit Access during the Recovery”. In: *Harvard Business School / Harvard Kennedy School*. HBS working paper; discusses informational opacity, collateral limits, and risk in small business lending.

Mourad, Frederico A., Rafael F. Schiozer, and Toni R. E. dos Santos (2020). *Bank Loan Forbearance: Evidence from a Million Restructured Loans*. Working Paper Series 541. Central Bank of Brazil, Research Department.

Office of the Comptroller of the Currency (Dec. 2014a). *Bankers' Guide to the SBA 7(a) Loan Guaranty Program*. Tech. rep. OCC.

— (2014b). *The Bankers' Guide to the SBA 7(a) Loan Guaranty Program*. Tech. rep. Describes the SBA 7(a) program as serving creditworthy small business borrowers who cannot otherwise obtain credit at reasonable terms and lack alternative financing sources. U.S. Department of the Treasury.

Office of the Comptroller of the Currency (2020). *Comptroller's Handbook: Small Business Lending*.

Pesendorfer, Martin and Philipp Schmidt-Dengler (2003). *Identification and Estimation of Dynamic Games*. NBER Working Paper 9726. NBER / LSE.

Petersen, Mitchell A. and Raghuram G. Rajan (1994). "The Benefits of Lending Relationships: Evidence from Small Business Data". In: *Journal of Finance* 49.1, pp. 3–37.

Pflüger, Michael P. and Jens Südekum (2013). "Subsidizing Firm Entry in Open Economies". In: *Journal of Public Economics* 97. IZA Discussion Paper No. 4384 (2009), pp. 258–271.

Roberts, Michael R. (2015). "The Role of Dynamic Renegotiation and Asymmetric Information in Financial Contracting". In: *Journal of Financial Economics* 116.1, pp. 61–81.

Rosenbaum, Paul R. and Donald B. Rubin (1983). "The Central Role of the Propensity Score in Observational Studies for Causal Effects". In: *Biometrika* 70.1, pp. 41–55.

SBA 7(a) Loans (2025). *What Is the SBA Express Loan?* https://www.sba7a.loans/sba-7a-loans-small-business-blog/what-is-the-sba-7a-express-loan/. Accessed: 2025-10-23.

SBA SOP (2023). See Subpart B, Ch. 2, Equity Injection. U.S. Small Business Administration.

Stachurski, John (2009). *Economic Dynamics: Theory and Computation*. Cambridge, MA: MIT Press.

Stillerman, David (2024). "Loan Guarantees and Incentives for Information Acquisition". SSRN Working Paper, posted April 17, 2023; last revised January 26, 2024.

Stohs, Mark H. and David C. Mauer (1996). "The Determinants of Corporate Debt Maturity Structure". In: *Journal of Business* 69.3, pp. 279–312.

Stokey, Nancy L., Jr. Robert E. Lucas, and Edward C. Prescott (1989). *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.

STR / CoStar Group (2024). *The Role of Total Revenue Management in a Hotel Profitability Strategy*. Accessed 2025.

Sufi, Amir (2009). "Bank Lines of Credit in Corporate Finance: An Empirical Analysis". In: *The Review of Financial Studies* 22.3, pp. 1057–1088.

Townsend, Robert M. (1979). "Optimal Contracts and Competitive Markets with Costly State Verification". In: *Journal of Economic Theory* 21.2, pp. 265–293.

U.S. Bureau of Economic Analysis (2025). *GDP by Industry: Value Added by Industry (Table TVA 105)*. Interactive Data application. URL: https://apps.bea.gov/iTable/?reqid=1603&step=2&Categories=GDPxInd&isURI=1#eyJhcHBpZCI6MTYwMywic3RlcHMiOlsxLDIsNF0sImRhdGEiOltbImNhdGVnb3JpZXMiLCJHRFB4SW5kIl0sWyJUYWJsZV9MaXN0IiwiVFZBMTA1Il1dfQ== (visited on 09/07/2025).

U.S. Bureau of Labor Statistics, Business Employment Dynamics (BLS BDM) (n.d.). *Establishment Age and Survival Data*. https://www.bls.gov/bdm/bdmage.htm. Accessed: 2025-10-13.

U.S. Department of the Treasury, Office of Financial Institutions Policy (2023). *Financing Small Business: Landscape and Policy Recommendations*. Tech. rep. U.S. Department of the Treasury.

U.S. Small Business Administration (2024a). *Small Business Administration Loan Program Performance*. Tables include 7(a) charge-off rates as a percent of unpaid principal balance by fiscal year. URL: https://www.sba.gov/document/report-small-business-administration-loan-program-performance.

U.S. Small Business Administration, Office of Advocacy (2024b). *Frequently Asked Questions About Small Business Finance, 2024*. Tech. rep. U.S. SBA.

U.S. Small Business Administration Office of Inspector General (2019). *Evaluation of SBA's 7(a) Loans to Poultry Farmers*. Report 18-13. U.S. Small Business Administration.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy (2008). "Markov Perfect Industry Dynamics With Many Firms". In: *Econometrica* 76.6, pp. 1375–1411. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA6158.

9 Appendix

9.1 Small Business Lending in Top 10 Economies

Table 7: Estimated Government-Backed SME Loan Assistance (2010–2019) in Top 10 Economies

Country	Primary program(s)	USD (bn)	Notes / sources
USA	SBA 7(a) Loan Program	~215.1	Sum of official annual totals.[a]
China	Inclusive finance to small and micro firms	~575	Outstanding balance proxy; uncertainty.[b]
Japan	Japan Finance Corporation (JFC) SME loans	~50	Aggregated from JFC annuals; SME share assumed stable.[c]
Germany	KfW domestic SME promotional loans	~420	Extrapolated from KfW financing volumes.[d]
UK	Enterprise Finance Guarantee; Start Up Loans	~15	Parliamentary and BBB publications.[e]
France	Oséo / Bpifrance guarantees and co-financing	~90	Oséo 2010 baseline, continued under Bpifrance.[f]
India	PMMY and related schemes	~175	PMMY sums (2015–2019) plus conservative pre-2015.[g]
Italy	Fondo di Garanzia per le PMI	~120	Back-cast from 2018–2019 volumes.[h]
Brazil	BNDES MSME credit lines	~180	MSME disbursement averages from BNDES.[i]
Canada	Canada Small Business Financing Program (CSBFP)	~8.5	Gov't CSBFP time series; year-by-year FX conversion.[j]
	Total	**~1600**	

[a] SBA newsroom releases.
[b] People's Bank of China inclusive-finance statistics; outstanding SME balance at decade end.
[c] JFC annual reports.
[d] KfW financial reports and publications.
[e] British Business Bank reports; UK Parliament EFG statistics.
[f] Bpifrance publications; Oséo 2010 annual report baseline.
[g] MUDRA/PMMY reports.
[h] Fondo di Garanzia data and reports.
[i] BNDES transparency downloads; MSME series.
[j] Government of Canada CSBFP reviews and statistical reports (2010–2019).

9.2 7(a) Details

9.2.1 Top Recipient Sectors: 2010-2019

	NaicsDescription	Count (k)	Amount (USD, B)	Share (counts)	Share (amount)
1	Hotels (except Casino Hotels) and Motels	11.99	24.45	0.02	0.08
2	Full-Service Restaurants	28.23	12.17	0.04	0.04
3	Gasoline Stations with Convenience Stores	8.418	9.067	0.01	0.03
4	Offices of Dentists	12.87	8.012	0.02	0.03
5	Limited-Service Restaurants	19.54	7.641	0.03	0.03
6	Child Day Care Services	9.020	6.825	0.01	0.02
7	Offices of Physicians (except Mental Health Specialists)	9.997	4.694	0.01	0.02
8	Car Washes	3.530	4.452	0.00	0.01
9	Veterinary Services	5.495	4.419	0.01	0.01
10	Fitness and Recreational Sports Centers	11.33	4.235	0.02	0.01
11	Beer, Wine, and Liquor Stores	7.768	4.106	0.01	0.01
12	Broilers and Other Meat Type	3.762	3.556	0.01	0.01
13	General Automotive Repair	10.35	3.422	0.01	0.01
14	All Other Specialty Trade Contractors	10.38	3.323	0.01	0.01
15	Supermarkets and Other Grocery (except Convenience) Stores	5.093	3.319	0.01	0.01

Notes: Count is displayed in thousands (k). Amount is displayed in billions of USD (B). Rounding uses up to 4 significant digits for values over 100 to keep at most 5 digits visible in each numeric field.

9.2.2 7(a) Loan Characteristics

We plot characteristics of the SBA 7(a) loans that were granted.

Figure 11 illustrates the distribution of approved loan amounts. The majority of loans are relatively small, ranging from $50k to $5 million, with an average loan amount of $383,556. This shows that SBA loans are designed to be accessible to small businesses that might not qualify for larger financing options.

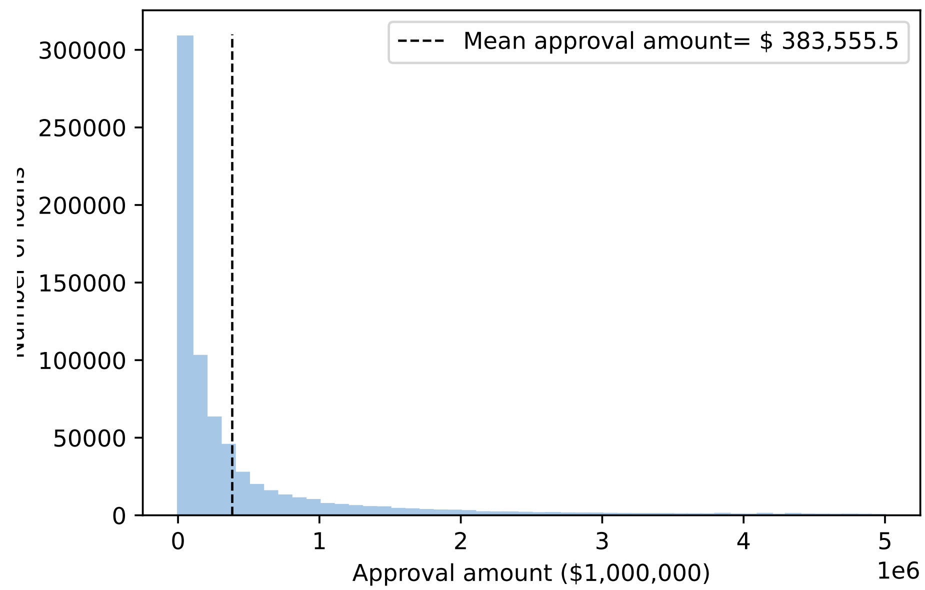

Figure 11: Approved loan amount (in million USD)

Loan terms are generous, as depicted in Figure 12. The average loan term length is 10 years, providing borrowers with extended periods to repay their loans, which can help manage cash flow and reduce the immediate financial burden on small businesses.

Another critical aspect of SBA loans is the guarantee provided by the SBA, shown in Figure 13. On average, SBA loans are heavily guaranteed at 66%, which mitigates the risk for lenders and makes it easier for small businesses to obtain financing.

However, the default rate on SBA loans is significant: our data shows the default rate is approximately 2-6% across years, highlighting the risk involved in lending to small businesses. This variability in default rates underscores the importance of the SBA guarantees in sustaining lender participation.

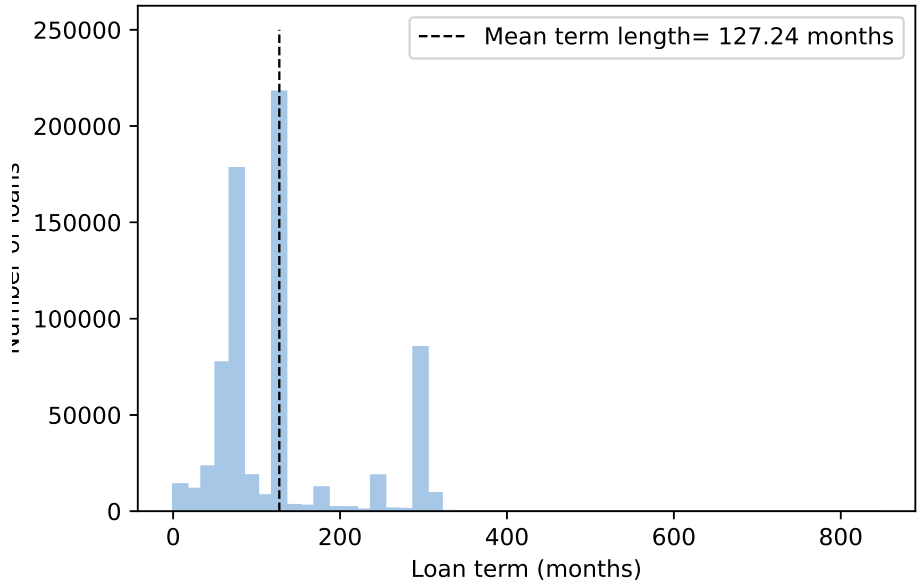

Figure 12: Term lengths are generous

7(a) margins

To assess whether the 7(a) program expands credit access, we briefly examine realized loan margins for banks. While not all loans have matured, we focus on the subset that has, using these outcomes to calculate rough accounting margins. For each loan l issued by bank b, we define the real interest earnings as:

$$\pi_b^l = r_{\text{real}}^l \cdot P_b^l \tag{14}$$

where r_{real}^l is the real (inflation-adjusted) interest rate and P_b^l is the principal.

Loan losses are given by charged-off amounts:

$$l_b = \sum_l P_{\text{chargeoff}}^l \tag{15}$$

We then compute an approximate realized margin for each bank:

$$m_b = \frac{\sum_l \pi_b^l - l_b}{\sum_l P_b^l} \tag{16}$$

Margins are near zero on average—and negative for many banks—when guarantees are excluded. Once SBA guarantees are included, margins rise significantly, with the

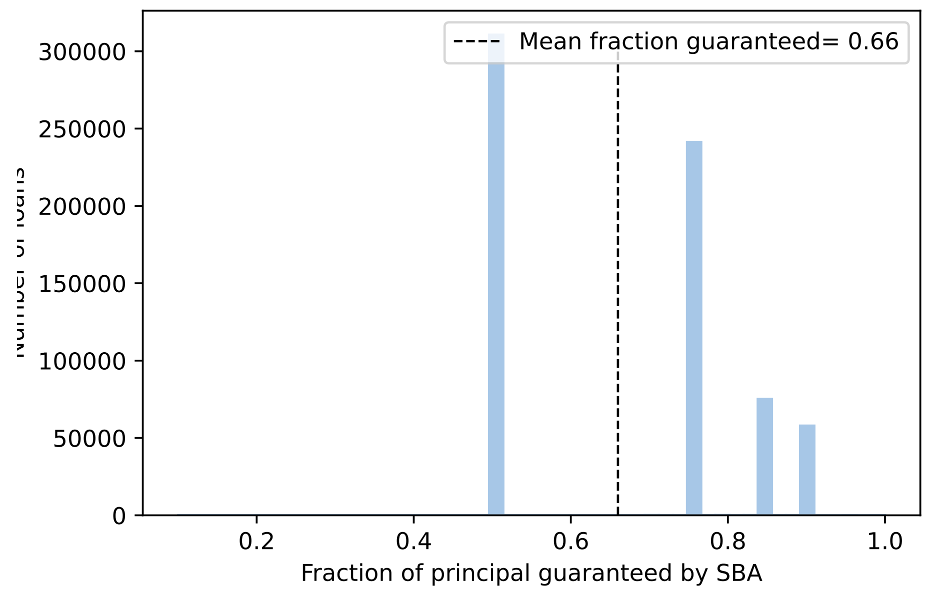

Figure 13: Loan guarantees are large

program-wide average increasing to roughly 4%. This suggests that, for a large share of recipients, 7(a) loans would not have been privately profitable absent the guarantee.

9.3 Data construction

9.3.1 Scraping hotels data

We construct a map of local competitors of hotel markets for our analysis.

We use google maps as a meta-aggregator of hotels and motels.An attractive feature using google for aggregating this information is that we can find many hotels and motels which are otherwise not available on commercial aggregators, since owners of such businesses can create their own location without needing to have a business relationship with the aggregator- google. Another attractive feature We query for 'hotels' in the maps search in every city that has atleast one hotel recipient of SBA loans.

A limitation with using this querying exercise at the city-level is that for large cities, hotel options are limited to a curated set of N options, where N is bounded by 60. This means that unless a query looks within a particular cell of a city by zooming into it, we would only scrape a small portion of the market. While technically possible to capture competitive effects with a subsample, we lose power and could be biased by selection into the set of hotels that maps curates.

We solve this by scraping at a fine level of geographic aggregation, corresponding to a particular 'zoom' within google maps. We scrape over many zoomed cells for large cities and stitch them together to create a dataset that corresponds to the city.

We use a commercial aggregator of scrapes—Apify—to operationalize this scraping exercise. We extract data from all these cities for over 5 waves, where each wave corresponds to blah blah. By parallelizing the scraping exercise, and randomizing over the order of scrape queries, we ensure no selection into data extraction and execute it within 5-6 hours.

A caveat is that in order to create boundaries of cities, our scraper uses *nominatim* boundaries. These are not always available or reliable, so we are unable to extract data for 421 cities. These are typically smaller cities.

9.3.2 Creating hotel markets

City sizes and populations in our hotel geographies range from 8 million to less than $10,000$.

We focus on the competition effects of hotels in various markets. Since most consumers select for hotels in a particular area (Mody et al. (2023)) we will operationalize this by breaking up the largest cities into smaller consumer relevant submarkets.

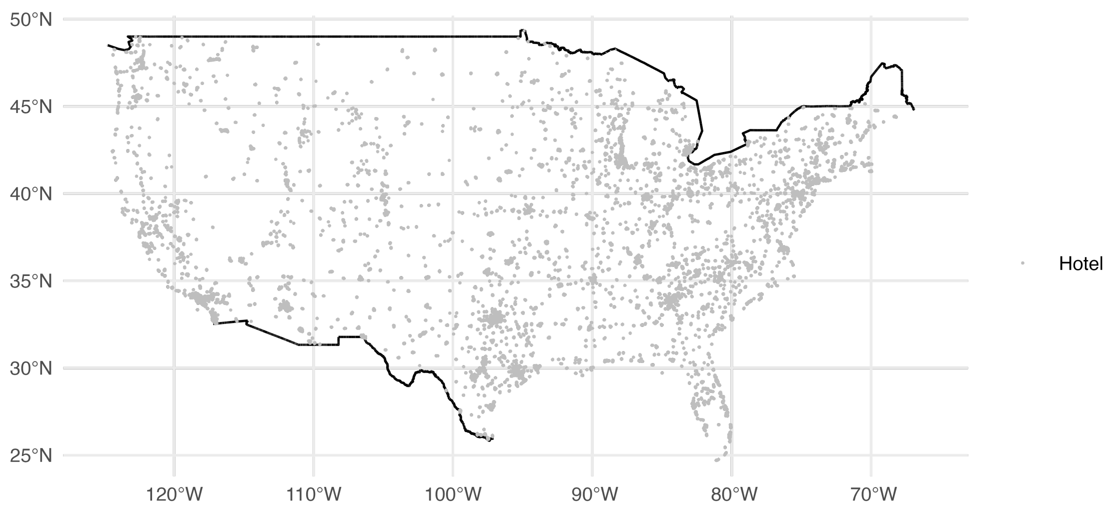

Figure 14: Dots represent hotel markets scraped in the 48 states from Google maps by stitching together sub-city level queries.

We select all cities larger than size $400,000$ to be broken up into smaller submarkets. We impose limits on how small these submarkets can be, the minimum size is $150,000$ and the maximum size is $650,000$

We accomplish this in three steps-

1. Break cities into 'islands' separated by water bodies.

2. Cluster zipcodes and neighborhoods within islands.

3. Aggregate small submarkets.

As the first step, we break up cities into smaller geographic regions when they are separated by a water body- a river for example.

Subsequently, we create markets by aggregating up smaller geographic areas. A challenge here is that there are no standardized geographic units between cities and zipcodes that are comparable across all cities in the US. So we proceed by constructing our own aggregates from zipcodes. A feature of our data that helps us in this exercise is that many

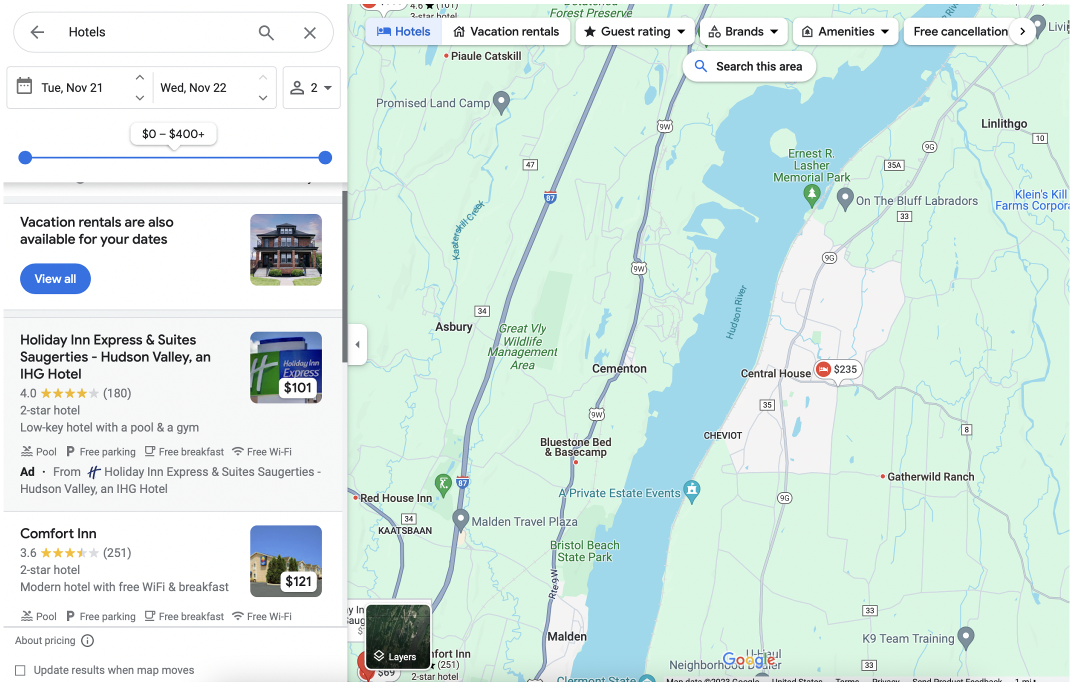

Figure 15: Sample town queried for its hotels on Google maps

zipcodes are scraped with neighborhood names, which act as natural pre-clusters for many zipcodes. These are helpful because they club together areas which are informally sub-unit of many cities and

We start by combining zipcodes which have neighborhood information together to give pre-clusters. We just use the individual zipcode as a pre-cluster for zipcodes for which we don't find neighborhood information from google. The aggregation is done in a way to target a certain number of markets depending on the city population (source-ACS 2019). We change this size from N=2 for cities of size 400,000 to N=10 for a city of size 10,000,000 people. We fit a logarithmic curve to scale the number of markets across cities.

We aggregate data up from these 'pre-clusters' by using k-means clustering algorithm. The objective function that this algorithm solves attempts to split the geographic data, based on pre-cluster mean latitude-longitude, into clusters of size n such that the distance between the centroid of the cluster and each individual unit within this cluster is minimised. A great advantage of this method is that because our pre-cluster units lie in a 2-dimensional geographic space, the clustering bins contiguous pre-clusters together into one unit. This is contrary to some other natural approaches which might seek to group units together based on other information or with constraints (see below for an alternative that doesn't work well). We end up with 2679 markets in our sample with

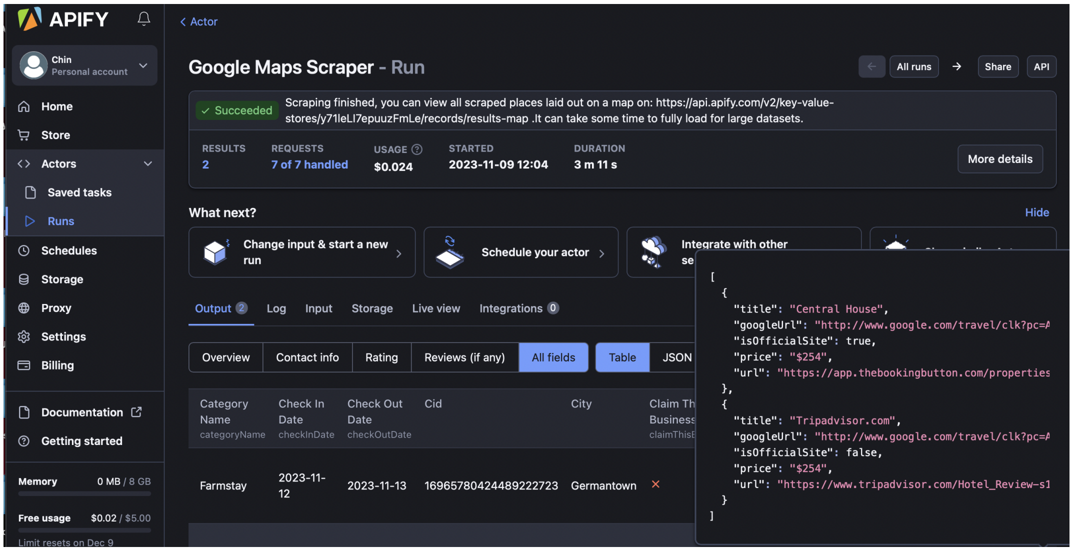

Figure 16: Query scraped to fetch information about prices and characteristics

population sizes.

We attempted two alternative ways to create these markets- 1) aggregating up directly from zipcodes 2) using weight-constrained k-means clustering. The former achieves contiguity, but since it doesn't incorporate the additional grouping provided by neighborhoods, it often splits zipcodes which should naturally be a part of the same market into two different clusters.

We tried using a weight constrained k-means clustering approach to balance dual objectives of having clusters which are together in space, with the added constraint that the markets we end up constructing are roughly of similar population size. This is an NP hard problem, for which a simplification exists using min-cut/max-flow algorithms. However, since zipcodes are very differently populated across space, the constrained optimisation routine often breaks contiguity by trying to achieve similar bundles of zipcodes in terms of population. We provide code for these two alternatives as well, since this could be promising in other settings where aggregation is across units of similar weight, but the final weight is something we want to balance.

We clean the gathered hotel data across cities to have a dataset which has hotel prices, address, geolocation (lat/long), star ratings distribution, hotel type.

While doing this exercise, we make use of the many characteristics available for each hotel from the google query. We extract rich additional information- review tags, amenities information, as well as information on 'nearby hotels'. We use the former two to construct rich covariates to predict gaps in hotel-star type and prices in our dataset. We

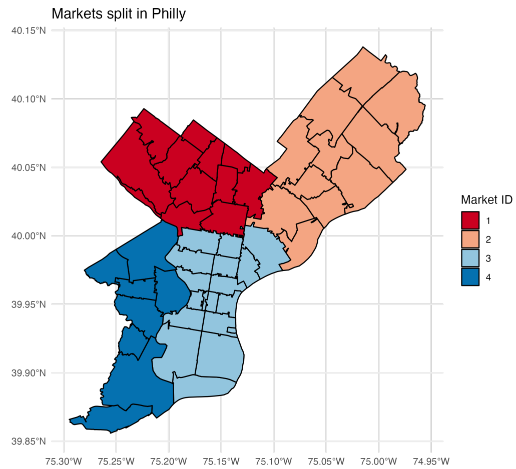

Figure 17: Markets in a sample city–Philadelphia–after the splitting exercise

use the latter to extend the markets that are constructed in our dataset.

9.3.3 Extracting rich covariates

We select additional information on hotels such as the availability of free Wi-Fi, free breakfast, free parking, outdoor pool, air-conditioned, pet-friendly etc. We also extract all review-tags, which are NLP extracts from text reviews left by google users for these hotels. These include keywords such as 'clean', 'comfortable', 'desk','sheets' which have additional information about hotels that are unobserved from the hotel listing but derived from customer experiences. We aggregate review tags from all hotels and find more than 8000 unique tags across hotels. We select the 200 most commonly tagged ones. Finally, we also extract information from hotel descriptions by bagging the most common words found in the hotel description. We tokenize the description, remove stop words and non-descriptors and select top 200 most common words found in the description. Then, we assign a true/false value to whether these words are present in the hotel description.

For most hotels, not all of these additional information is available so instead of using this information directly as a rich, but sparse set of covariates, we use it to impute missing information as described below.

9.3.4 Adding nearby hotels

Google records nearby hotels which are relevant to each hotel queried in our dataset. It also records prices, address, description, star rating, reviews for each of these. We use this additional information to extend our market definitions and include additional hotels which are present in the market. We do this, by selecting the hotels in this nearby list that have not already been scraped, scraping latitude and longitude information from google and subsetting to hotels which have price information available.

We extract additional information about these hotels by bagging the most common words found in the hotel description. We tokenize the description, remove stop words and non-descriptors and select top 200 most common words found in the description. Then, we assign a true/false value to whether these words are present in the hotel description. We will use these to construct a rich classifier for hotel star categories for some of these nearby hotels.

9.3.5 Classifying hotel types and identifying franchises

We have a data-field called hotel category available to us from the data collection exercise. We use this to determine whether an accommodation is of the type 'hotel', 'motel' or 'other', which is a catch all for campgrounds, lodges etc. Since businesses often carry multiple category tags, we use a hierarchical approach- if a hotel carries a category tag motel, it's classified as a motel regardless of whether it has other tags. If the listing carries no motel tag, but carries a hotel tag, we classify it as a hotel. If neither of these are true, we classify it as other.

Hotel names in our data carry brand information for franchises. We use this to characterize the ownership structure of hotels in our dataset. We first collect a list of franchise names derived from (a) SBA loans dataset, which lists franchises, and (b) manually identifying the most common parent companies and their sub-brands from various internet sources such as STR demos, Statista and Wikipedia. After constructing a dictionary consisting of 160 main brands (eg. Hilton) and sub-brands (eg. Doubletree), we merge them with our scraped dataset based on whether the hotel name or description carries franchise name. Close to a half of hotels in our dataset are a part of a franchise.

9.3.6 Filling in hotel stars and prices

We impute the values of hotel stars and prices for hotels that are missing this information from our dataset. Nearly 25% of our main scraped data lack price and hotel-star information and nearly 40% of our nearby hotels data lack hotel-star information (note that we already restricted ourselves to the hotels here that have prices). Using the rich information collected from review tags, additional amenities information and NLP processing from descriptions, we predict the hotel stars and prices.

We use Gradient Boosted decision trees, and XGBoost algorithms. These are ensemble methods, that use a large number of non-linear classifiers and predictors respectively to solve a non-linear supervised learning problem.

We predict the hotel star category for the 40% of 'nearby' hotels data for which the star category is not mentioned explicitly in the hotel description. We use the remaining data as a training set and train a gradient boosted decision tree model using 5-fold validation of parameters tuned over a 3 dimensional grid on- size of the ensemble, learning rate and depth of the trees. We achieve an in-sample classification rate of 75% accuracy with no predicted star rating more than 1 star off. We then predict the hotel stars for the rest of the sample.

Some hotels lack prices, and some hotels lack star ratings for the main scraped dataset. We first predict the hotel stars using all these characteristics, using a similar approach of training, validation and prediction as for the nearby hotels. Then, we create 2 additional datafields- the mean price for a given zipcode and whether or not the hotel stars have been imputed or not. We use these covariates to train a regression tree model using XGBoost. We tune parameters using a 5-fold cross validation on a grid search over a 6 dimensional grid- size of the ensemble, learning rate, depth of the trees, minimum child weight for a tree, subsample size used in model and subsample size used in individual regressor. We train the model to predict the ratio of hotel price to the average price. This procedure yields substantial gains over predicting price directly or imputing it simply using local averages since this procedure incorporates both fundamental differences in underlying local costs, by taking mean prices, as well as accounting for premia associated with having other characteristics of the hotel throughout the sample.

These predicted prices are used to construct differentiation instruments.

9.3.7 Collecting Hotel Tax Data

9.4 Local Lodging Tax Data: Sources and Construction

The dataset comes from four inputs: *HVS Lodging Tax Report, USA* (scope, definitions), *HVS Lodging Tax Report USA, with citytaxes* (component rates by urban center), *Avalara Tax Rate Explanation* (Avalara guidance on ZIP assignment risk, obtained from firm), and an *Airbnb Lodging Tax Data* file that marks platform collection coverage.

HVS provides the backbone; Avalara informs how ZIP boundaries can misstate rates; the Airbnb file tags short-term rental collection.

Construction.

1. *Compile and normalize HVS.* We parse the city tables into a tidy structure with component ad valorem rates: state, county, city, and special district, the HVS published total, and a notes field. Then, we standardize identifiers (state postal, canonical city string) and put all rates on a common percentage basis. When multiple special districts apply inside a city footprint, we adopt the urban-center convention and use the *highest* applicable special-district rate for totals. Finally, we create the ad valorem subtotal

   $$\texttt{total\_advalorem} = \texttt{state} + \texttt{county} + \texttt{city} + \texttt{special}.$$

   We then compare `total_advalorem` to the HVS published total. If the absolute gap exceeds **0.05 percentage points**, we flag a reconciliation issue.

2. *Handle mixed regimes (excises, minimums, tiers).* Many jurisdictions add fixed per-night excises or minimum charges. We record any dollar fee as `excise_per_night`. For comparability, we convert to an effective percentage using a reference ADR of **$100** unless a source gives a city-specific ADR:

   $$\texttt{excise\_eff\_pct} \; = \; 100 \times \frac{\texttt{excise\_per\_night}}{\texttt{ADR\_ref}}, \quad \text{with } \texttt{ADR\_ref} = 100.$$

   For minimum taxes expressed as a dollar floor per room night, we treat the floor as an excise at `ADR_ref` and compute the effective add-on the same way. For tiered schedules (rate depends on price), we evaluate the tier at `ADR_ref` and retain the raw tier text in `notes`.

   $$\texttt{total\_effective} \; = \; \texttt{total\_advalorem} + \texttt{excise\_eff\_pct} \quad \text{and} \quad \texttt{total\_advalorem}.$$

3. *Enrich and QC.* At the end, we attach an `airbnb_collects` tag in {`none`, `state`, `local`, `both`} from *airbnb_lodgingtax* at the most granular match available. We avoid ZIP-only assignments when joining to other assets. When a downstream join must use ZIPs, set `zip_risk='elevated'` if a ZIP maps to more than one county or more than one city in the ZIP master, and `zip_risk='low'` otherwise.

4. *Cleanup* We apply targeted state adjustments where statewide frameworks differ in structure (for example, Hawaii transient accommodations tax layered with general excise, Louisiana parish and city overlays) and document each override in a `qc_status` note.

Output. The released table has one row per city with identifiers (`city`, `state_postal`), component rates (`state_rate`, `county_rate`, `city_rate`, `special_rate`), totals (`total_advalorem`, `excise_per_night`, `excise_eff_pct`, `total_effective`), short-term rental coverage (`airbnb_collects`), provenance (`source_doc`, `source_page`, `as_of`), and QC indicators (`max_special_applied`, `recon_flag`, `zip_risk`, `qc_status`). This structure keeps the HVS methodology intact, makes mixed regimes comparable on a common scale.
,These tax rates are used to construct effective consumer prices. Since they are common to all hotels, they act as effective shifters of the value of the inside nest options, giving us plausibly exogenous variation to identify the nest correlation parameter.

9.5 SBA's Role in Hotel Financing

The hotel financing structure illustrates the credit gap that 7(a) loans are designed to address. Consider a typical 50-room limited-service hotel with approximately \$6 million in total project costs. Conventional lenders finance only hard assets (land and building) at typical 75% loan-to-value ratios. For a typical \$6 million hotel project, this yields roughly \$3.15 million in debt capacity. However, hotels require substantial additional capital for furniture, fixtures, equipment, pre-opening marketing, inventory, franchise fees, and working capital reserves—collectively representing \$1.5-2 million in non-collateralizable expenses. With owner equity of 15-25%, a financing gap of \$1-2 million emerges. (Hospitality Financial and Technology Professionals 2024).

With typical owner equity of 15-25% (\$0.9-1.5M), a financing gap of \$1-2 million emerges. The SBA 7(a) program addresses this gap through its statutory design: the 'credit elsewhere' test requires recipients cannot obtain credit on reasonable terms elsewhere, while SBA policy prohibits lenders from declining loans solely due to insufficient collateral. This enables 7(a) loans to provide mixed-purpose financing combining real estate acquisition with working capital and other expenses.

9.6 Summary Statistics of markets

	N	Mean	Std	Min	Median	Max
Price	48,779	124.99	104.13	2.0	100.90	9500.0
Hotel Stars	48,779	2.47	0.67	1.0	2.00	5.0
Number of reviews	48,779	568.53	1986.42	0.0	309.00	129009.0
Average rating	47,288	3.94	0.63	1.0	4.06	5.0

Table 8: Summary statistics for hotels in the scraped dataset

9.7 Two-Part Debt Structure

Hotels face two distinct financing needs: property acquisition and working capital. These differ fundamentally in collateralizability and lender risk. We capture this through a two-part debt structure.

Fully Collateralized Mortgage Debt. The first model component finances property acquisition. This debt is fully collateralized and poses no risk to lenders, maintaining a risk-free price in all states of the world. We model mortgage payments as proportional to expected profits: $m_{jt} = a\bar{\pi}_{jt}$, where $\bar{\pi}_{jt}$ denotes the ex-ante expected profit of firm j at time t; this object is defined more precisely in the next section.

This proportional specification reflects size-based scaling: larger hotels generate higher revenues and carry larger property investments. We use expected profits rather than revenue as our size proxy because profits in the model scale closely with revenue, while providing a more natural state variable for the firm's dynamic problem. Industry metrics confirm that revenue, investment, and profits all scale with physical attributes such as room count and square footage.[49]

This structure reflects the long-term, low-risk nature of real estate financing in the hotel industry. Its main effect is to reduce the residual cash flow available to the firm each period. As a, the share of average profit pledged to mortgage payments, increases, so does the effective volatility of the remaining profit stream. [50]

[49]See Cushman & Wakefield 2018, American Hotel & Lodging Association 2025

[50]Commercial real estate (CRE) loans—such as those used to finance hotel property—are typically long-term, collateralized instruments with repayment horizons of 10 to 30 years. Many are securitized as commercial mortgage-backed securities (CMBS), which are considered low-risk in senior tranches due to high recovery rates and structured servicing. See Trepp (2023), "U.S. Hotel Sector: CMBS Performance and Credit Risk Trends"; Mortgage Bankers Association (2021), "Commercial/Multifamily Real Estate Finance: Facts and Figures"; and the SBA 504 Loan Program, which provides fixed-rate, long-term financing for commercial property acquisition.

Partially Recoverable Working Capital Loans. The second debt component finances remaining property expenses and working capital expenses—furniture, equipment, inventory, and day-to-day operations—that are only partially recoverable in the event of default. We model this as a loan with a constant recovery rate $\omega < 1$, calibrated to empirical averages from SBA loan data. Unlike mortgage debt, this loan's terms vary: its interest rate increases with firm's default risk, which in turn depends on ω, the firm's leverage, and profit volatility.

As firms become more leveraged, credit becomes more expensive and eventually unavailable, forcing exit when losses occur. This margin is the key friction in our model and the central focus of the policy intervention.

Interpretation and Simplification. This two-part structure is stylized but empirically grounded. Hotels commonly maintain separate loans for property and operations, with distinct collateral structures (senior and junior liens) that prevent ex-post blending without restructuring.[51] Our setup nests a single partially secured loan (with $m_j = 0$) as a special case. The key modeling assumption is that the recovery rate ω is not under the firm's control.

By loading all risk into working capital debt, we capture the economically relevant distinction: property generates credible collateral while operational expenses do not. This is where "last mile" lending friction—and policy interventions like 7(a)—matter most. Mortgage debt will enter solely as a per-period cost, while working capital loans drive all variation in credit pricing and constraints.

We now introduce the firm's profit stream in abstract terms, highlighting the sources of uncertainty. We will later specify the details of accommodation-side mechanics that generate it.

9.8 Model Proofs

Natural borrowing bound. Assume:

(a) **Lender participation (LPC).** For any one–period loan (L_{jt}, r_{jt}) issued at t,

$$\mathbb{E}_t\left[(1 + r_{jt})L_{jt}\right] \;\geq\; (1 + m)\, L_{jt}.$$

[51]See Lodging Staff 2022 for a summary of how hotel lenders structure multiple loan types, including bridge, mezzanine, and construction financing. Legal structures such as senior and junior liens restrict collateral to specific obligations, making it infeasible to blend different loans ex-post without complete restructuring.

(b) **Refinancing feasibility.** At $t+1$ the borrower can always avoid default by raising new one–period debt; we let it choose the *minimal* amount needed (never borrow strictly more than the shortfall).

(c) **No–Ponzi (TVC).** $\lim_{T\to\infty} \mathbb{E}_t \left[\frac{L_{j,t+T}}{(1+m)^T} \right] = 0.$

Then, for any t, the outstanding principal satisfies the PV bound

$$L_{jt} \;\le\; \mathbb{E}_t \left[\sum_{s=1}^{\infty} \frac{\pi_{j,t+s}}{(1+m)^s} \right] \;\equiv\; \overline{L}_{jt}^{\,\mathrm{nat}}. \tag{17}$$

Proof. Refinancing feasibility with *minimal* new borrowing implies the one–step accounting inequality

$$(1+r_{jt})\, L_{jt} \;\le\; \pi_{j,t+1} + L_{j,t+1} \tag{18}$$

(equality when $\pi_{j,t+1} < (1+r_{jt})L_{jt}$; otherwise $L_{j,t+1} = 0$ and the inequality is slack). Taking $\mathbb{E}_t[\cdot]$ and using LPC,

$$(1+m)L_{jt} \;\le\; \mathbb{E}_t\big[(1+r_{jt})L_{jt}\big] \;\le\; \mathbb{E}_t\big[\pi_{j,t+1} + L_{j,t+1}\big].$$

Divide by $(1+m)$:

$$L_{jt} \;\le\; \mathbb{E}_t \left[\frac{\pi_{j,t+1}}{1+m} + \frac{L_{j,t+1}}{1+m} \right].$$

Iterating T times yields

$$L_{jt} \;\le\; \mathbb{E}_t \left[\sum_{s=1}^{T} \frac{\pi_{j,t+s}}{(1+m)^s} \;+\; \frac{L_{j,t+T}}{(1+m)^T} \right].$$

Let $T \to \infty$ and apply TVC to get (17). □

9.8.1 Debt Limit: Existence and Uniqueness

Proposition 1. *A finite period t debt limit $L_{jt}^{\max}$ exists and is unique for a firm with finite next-period debt limits $L_{jt+1}^{\max}$, provided the distribution of composite uncertainty $\tilde{\pi}_{jt} = \pi_{jt}^{\phi} + \pi_{jt}^{\xi} - \pi_{jt}^{C}$ is thin-tailed in the following weak sense:*

$$\lim_{x\to\infty} \frac{(1 - F_{\tilde{\pi}}(x))^2}{x f_{\tilde{\pi}}(x)} = 0,$$

where $F_{\tilde{\pi}}(x)$ and $f_{\tilde{\pi}}(x)$ denote the CDF and PDF of $\tilde{\pi}_{jt+1}$, respectively.

Fix a next-period loan cap $L''_{\max} \in (0, \infty)$. Let $L' > 0$ be the current loan (maturing next period) and $r' \geq 0$ the interest rate. Let ν be uncertainty with CDF F, PDF f, survival $S := 1 - F$. This can be interpreted as being the mixture of uncertainty over types, demand and cost, as they are written additively in the model. Given (L', r'), default occurs when

$$\nu \leq \bar{\nu}(L', r') \; := \; (1 + r')L' - \pi' - L''_{\max},$$

so $\Pr[\text{default}] = F(\bar{\nu})$ and $\Pr[\text{no default}] = S(\bar{\nu})$. If default occurs, the bank recovers a fraction $\omega \in [0, 1]$ of face value $(1 + r')L'$. Write the bank's *total* expected profit as $\mathbb{E}\left[\pi^{\text{bank}}(r', L', \pi')\right]$.

It is convenient to work with *per-unit* profit:

$$\Pi(L', r') \; := \; \frac{\mathbb{E}\left[\pi^{\text{bank}}(r', L', \pi')\right]}{L'} \; = \; r' + F(\bar{\nu})(\omega - 1 - r') \; = \; S(\bar{\nu})\, r' + F(\bar{\nu})\,(\omega - 1). \tag{1}$$

The dividend requirement

$$\mathbb{E}\Big[\pi^{\text{bank}}(r^*, L', \pi')\Big] \Big|_{L' = L_j'^{\max}} \; = \; mL'$$

is equivalently

$$\Pi(L', r'^*(L')) \; = \; m, \tag{2}$$

with $m \geq 0$.

Profit-maximizing rate and identity. Assuming an interior solution, the FOC $\partial \Pi / \partial r' = 0$ and $\partial \bar{\nu} / \partial r' = L'$ yield

$$0 \; = \; S(\bar{\nu}) + (\omega - 1 - r')\, f(\bar{\nu})\, L' \quad \Longrightarrow \quad r'^*(L') \; = \; (\omega - 1) + \frac{S(\bar{\nu})}{L' f(\bar{\nu})}. \tag{3}$$

Substituting (3) into (1) gives the *maximal per-unit profit*:

$$\Pi\big(L', r'^*(L')\big) \; = \; (\omega - 1) + \frac{S(\bar{\nu})^2}{L'\, f(\bar{\nu})} \, . \tag{4}$$

Hence the dividend requirement (2) is equivalent to

$$\omega - 1 \; + \; \frac{[1 - F(\bar{\nu})]^2}{L'\, f(\bar{\nu})} \; = \; m \, . \tag{5}$$

Tail condition. Assume the mild tail condition

$$\lim_{x \to \infty} \frac{S(x)^2}{x\, f(x)} \; = \; 0. \tag{A}$$

This holds for thin-tailed (e.g. Normal, Exponential), Lognormal, and Pareto($\alpha > 0$) distributions.

Existence of a finite current debt limit.

Proposition 2. *Fix $L''_{\max} < \infty$ and suppose* (A) *holds. Then there exists $\overline{L}' < \infty$ such that for all $L' > \overline{L}'$, $\max_{r' \geq 0} \Pi(L', r') < m$. Equivalently, the admissible set $\mathcal{A} = \{L' > 0 : \exists r' \geq 0 \text{ with } \Pi(L', r') \geq m\}$ is bounded above, and*

$$L_j'^{\max} \;:=\; \sup \mathcal{A} \;<\; \infty.$$

Proof. For any $r' \geq 0$, $\bar{v}(L', r') = (1 + r')L' - \pi' - L''_{\max} \geq L' - (\pi' + L''_{\max})$, so $\bar{v}(L', r') \to \infty$ at least linearly in L'. By (4),

$$\Pi\big(L', r'^*(L')\big) \;=\; (\omega - 1) + \frac{S(\bar{v})^2}{L' f(\bar{v})}.$$

Let $x := \bar{v}(L', r'^*(L'))$. Then $x \to \infty$ as $L' \to \infty$, and by *(A)*

$$\frac{S(\bar{v})^2}{L' f(\bar{v})} \;=\; \frac{S(x)^2}{x f(x)} \cdot \frac{x}{L'} \;\longrightarrow\; 0 \cdot (\text{finite}) \;=\; 0.$$

Hence $\lim_{L' \to \infty} \Pi(L', r'^*(L')) = \omega - 1 \leq 0$. If $m > 0$, there exists $\overline{L}'$ such that for all $L' > \overline{L}'$, $\Pi(L', r'^*(L')) < m$, and thus $\max_{r' \geq 0} \Pi(L', r') < m$. If the interior FOC fails so that the optimum is the corner $r' = 0$,

$$\Pi(L', 0) \;=\; F(\bar{v})\,(\omega - 1) \;\xrightarrow[L' \to \infty]{}\; \omega - 1 \;\leq\; 0,$$

and the same conclusion follows. Therefore $L_j'^{\max} < \infty$. $\square$

9.8.2 Debt Limit: Comparative Statics

Define the maximized per-unit value

$$V(L'; \pi', L''_{\max}) \;:=\; \sup_{r' \geq 0} \; \Pi(L', r'; \pi', L''_{\max}). \tag{CS.3}$$

The (current-period) feasible set of loans is

$$\mathcal{A}(\pi', L''_{\max}, m) \;:=\; \big\{\, L' > 0 : \; V(L'; \pi', L''_{\max}) \geq m \,\big\},$$

and the credit limit is

$$L'_{\max}(\pi', L''_{\max}, m) \;:=\; \sup\, \mathcal{A}(\pi', L''_{\max}, m).$$

Key monotonicity of the objective. Recall $\bar{v} = (1+r')L' - \pi' - L''_{\max}$,

$$\frac{\partial \Pi}{\partial \bar{v}} = S'(\bar{v})\, r' + F'(\bar{v})\,(\omega - 1) = -f(\bar{v})\, r' + f(\bar{v})\,(\omega - 1) = f(\bar{v})\big(\omega - 1 - r'\big) \;\le\; 0, \quad \text{(CS.4)}$$

since $r' \ge 0$ and $\omega - 1 \le 0$. Therefore,

$$\frac{\partial \Pi}{\partial \pi'} = \frac{\partial \Pi}{\partial \bar{v}} \cdot \frac{\partial \bar{v}}{\partial \pi'} = \frac{\partial \Pi}{\partial \bar{v}} \cdot (-1) \;\ge\; 0, \qquad \frac{\partial \Pi}{\partial L''_{\max}} = \frac{\partial \Pi}{\partial \bar{v}} \cdot \frac{\partial \bar{v}}{\partial L''_{\max}} = \frac{\partial \Pi}{\partial \bar{v}} \cdot (-1) \;\ge\; 0. \quad \text{(CS.5)}$$

Hence, for each fixed L' and any r', Π is (weakly) increasing in π' and $L''_{\max}$. Taking suprema over r' preserves monotonicity:

$$V(L'; \pi', L''_{\max}) \text{ is (weakly) increasing in } \pi' \text{ and in } L''_{\max}. \quad \text{(CS.6)}$$

Comparative statics of the debt limit.

Proposition 3 (Monotone comparative statics of $L'_{\max}$)**.** *Fix $\omega \in [0,1]$ and $m \ge 0$. Then:*

1. *$L'_{\max}(\pi', L''_{\max}, m)$ is (weakly) increasing in π' (higher revenue).*
2. *$L'_{\max}(\pi', L''_{\max}, m)$ is (weakly) increasing in $L''_{\max}$ (looser next-period cap).*
3. *$L'_{\max}(\pi', L''_{\max}, m)$ is (weakly) decreasing in m (tighter dividend requirement).*

Proof. (1) and (2): By (CS.6), for any fixed L', $V(L'; \pi'_2, L''_{\max}) \ge V(L'; \pi'_1, L''_{\max})$ whenever $\pi'_2 \ge \pi'_1$, and similarly for $L''_{\max}$. Hence the superlevel set $\{L' : V(L'; \cdot) \ge m\}$ expands as π' or $L''_{\max}$ increase. Therefore its supremum $L'_{\max}$ (weakly) increases.

(3): For any fixed $(\pi', L''_{\max})$, the feasible set $\mathcal{A}(\pi', L''_{\max}, m) = \{L' : V(L'; \pi', L''_{\max}) \ge m\}$ shrinks as m increases, so its supremum (weakly) decreases. □

Remarks.

- The argument uses only the sign in (CS.4), i.e. $r' \ge 0$ and $\omega \le 1$. No additional shape restriction on F is needed for these monotone conclusions.
- Identical reasoning shows $L'_{\max}$ increases in the recovery rate ω, since $\partial \Pi / \partial \omega = F(\bar{v}) \ge 0$.

9.8.3 The Space of 1-Period Debt Contracts

Borrowers fund their expenses through successive 1-period debt contracts. At entry (period 0) or after each period's profit is realized, the firm matches with a lender and contracts for debt L_{jt} at interest rate r_{jt}.

In their full generality, lending relationships an replicate arbitrarily complex repayment schemes. Formally, for any credit profile $c_{jt} = (\bar{\pi}_{jt}, \mathcal{H}_t)$, a lender i proposes a lending relationship by offering successive one-period contracts with history-dependent terms:

$$\mathcal{M}_j^i = \{\{L_{j\tau}^i(c_{j\tau}), r_{j\tau}^i(c_{j\tau})\}_{\mathcal{H}_\tau} \text{ for every history } \mathcal{H}_\tau)\} \tag{19}$$

To rule out Ponzi schemes, we impose a standard boundary condition on debt: $\lim_{T\to\infty} \mathbb{E}_t\left[\frac{L^i_{j,t+T}}{(1+m)^T}\right]$ 0. This condition implies finite debt limits at every history: lenders can only offer loans up to a maximum value $L_t^{\max}(c_{jt})$, equal to the firm's expected value in gross terms.

This general contract space is sufficiently rich to accommodate rigid amortization schedules and flexible refinancing arrangements. However, it poses tractability challenges. Because contract terms depend on the full profit history, the state space grows exponentially over time, obscuring the link between debt prices and economic fundamentals such as leverage, recovery rates, and profit volatility.

We make progress through two simplifying assumptions. First, lenders are identical and compete perfectly, ensuring that all loans earn exactly $(1 + m)$ in expectation. Second, borrowers can refinance any one-period loan costlessly with any lender. This eliminates relationship-specific rents and collapses the general contract space into a simpler structure.

Formally, after period t profit is realized, all lenders simultaneously propose 1-period contracts based on the borrower's credit profile c_{jt}. The borrower can freely match with any lender to obtain new debt L_{jt} and use this new debt to settle the previous obligation B^i_{jt}.

These assumptions place incumbent lenders on equal footing with competitors. If any lending relationship offers a continuation value exceeding $(1 + m)$ per dollar lent, a rival can undercut it by refinancing the borrower at a slightly better rate. In any renegotiation-proof competitive equilibrium, we obtain two key results:

- *Identity irrelevance:* Loan outcomes depend only on (c_{jt}, B_{jt}), not on lender identity.

- *Spot Equivalence:* Although lenders may offer full menus $\mathcal{M}_j$, only the current one-period term is enforceable. Every long-term debt contract reduces to pricing the current one-period loan or 'Spot' loan.

The second property is crucial. Although the lending environment is dynamic—histories evolve and continuation values matter—competition ensures that all rents are dissipated. Equilibrium pricing reduces to a sequence of static, per-period zero-profit conditions.

Lemma 1 (Nodewise zero profit)**.** *Fix any credit profile c_{jt} and accepted one–period loan (L_{jt}, r_{jt}) as a part of a lending relationship. Let p_{jt} denote the borrower's default probability. Under perfect competition and frictionless refinancing, the loan must break even at every node $\mathcal{H}_t$:*

$$(1 - p_{jt})(1 + r_{jt}) + p_{jt}\,\omega = 1 + m. \tag{20}$$

Proof. If an incumbent lender earns more than $(1 + m)$ on (L_{jt}, r_{jt}), a competitor can profitably refinance the loan at a slightly lower rate $r_{jt} - \varepsilon$, attracting the borrower while still earning $(1 + m)$. This deviation contradicts equilibrium. Conversely, if a lender earns less than $(1 + m)$ in the current period, they must expect to earn more than $(1 + m)$ in the future to break even overall—but the same argument rules out future rents. $\square$

This result implies that although lenders price debt based on future default risk and continuation values, equilibrium competition dissipates all surplus, period by period. With frictionless refinancing, pricing under dynamic lending is equivalent to a sequence of per-period zero-profit conditions. This "spot equivalence" mirrors similar results in sovereign debt models (e.g., Bulow and Rogoff 1989, Kehoe and Levine 1993, Kocherlakota 1996) and allows us to solve for equilibrium prices and debt limits recursively.

Importantly, the full dynamic structure remains intact: firms anticipate future prices when making entry and exit decisions, and lenders price today's loan based on tomorrow's state. Spot equivalence simply removes the need to track relationship-specific contracts or histories.

9.8.4 Existence and Uniqueness of the Value Function

Bellman operator in our model. Let the state be $s = (\bar{\pi}, L, \mathcal{H}, S) \in \mathcal{S}$. Fix S in this discussion. With discount $\beta \in (0, 1)$, define

$$(TV)(s) \;=\; \max\Big\{\, D(s) \;+\; \beta\,\mathbb{E}\big[V(S') \mid s\big] \;,\; -\,\omega\, L(s) \,\Big\},$$

where $0 \le L(s) \le B^{\max} < \infty$, $\omega > 0$, and $S' \sim Q(\cdot \mid s)$ is the Markov transition.

Theorem 1 (Existence and uniqueness in our setting)**.** *Suppose:*

(i) One–period payoff satisfies a linear growth bound in a shock c: there exist κ_0, κ_1 with $|D(s)| \le \kappa_0 + \kappa_1 |c(s)|$.

(ii) The shock follows a stable AR(1): $c' = \rho\, c + \varepsilon$ with $|\rho| < 1$, $\mathbb{E}[\varepsilon] = \mu_\varepsilon$, $\mathbb{E}[\varepsilon^2] = \sigma_\varepsilon^2 < \infty$.

(iii) Default penalty is bounded: $0 \le L(s) \le B^{\max}$.

Then the Bellman equation $V = TV$ admits a unique fixed point V^* in a weighted sup-norm space, and value iteration $V_{n+1} = TV_n$ converges to V^*.

We will show this by using a lemma for bounding the continuation value for the given problem.

Lemma 2 (Weighted Blackwell contraction under drift)**.** *Let $w : \mathcal{S} \to [1, \infty)$ and $B_w := \{V : \|V\|_w < \infty\}$ with $\|V\|_w := \sup_s |V(s)|/w(s)$. Assume:*

(A1) $|D(s)| \le C_D\, w(s)$ for some $C_D < \infty$.

(A2) (Drift) $\mathbb{E}[w(S') \mid s] \le \alpha\, w(s) + b$ with $\alpha \in [0, 1)$, $b < \infty$.

(A3) $0 \le L(s) \le B^{\max}$ so $|-\omega L(s)| \le \omega B^{\max}$.

Then $T : B_w \to B_w$ and

$$\|TV - TW\|_w \;\le\; \beta\alpha\, \|V - W\|_w \qquad \forall\, V, W \in B_w.$$

Hence T is a contraction (modulus $\beta\alpha < 1$) with a unique fixed point $V^* \in B_w$, and $V_{n+1} = TV_n$ converges to V^*.

Proof. Boundedness. For any $V \in B_w$,

$$|(TV)(s)| \le \max\{\, |D(s)| + \beta\, \mathbb{E}[|V(S')| \mid s],\, \omega B^{\max}\} \le C_D w(s) + \beta \|V\|_w\, (\alpha w(s) + b) \,\vee\, \omega B^{\max},$$

so $\|TV\|_w \le C_D + \beta\alpha \|V\|_w + \beta b < \infty$.

Monotonicity. If $V \le W$, then $\mathbb{E}[V(S') \mid s] \le \mathbb{E}[W(S') \mid s]$, and the outer $\max\{\cdot, -\omega L(s)\}$ preserves order, so $TV \le TW$.

Discounting. For any s,

$$\frac{|(TV)(s) - (TW)(s)|}{w(s)} \le \beta\, \frac{\mathbb{E}[\, |V(S') - W(S')| \,\mid s]}{w(s)} \le \beta\, \|V - W\|_w\, \frac{\mathbb{E}[w(S') \mid s]}{w(s)} \le \beta\alpha\, \|V - W\|_w.$$

Taking sup over s yields the claim. $\square$

Now, we can prove our original claim.

Corollary 1 (Verification for our model). *Under the primitives in Theorem 1, conditions (A1)–(A3) of Lemma 2 hold with $w(s) = 1 + c^2$. Hence V^* exists, is unique, and is obtained by value iteration.*

Proof (via Lemma 2 and verification). Apply Lemma 2 with the quadratic weight $w(s) := 1 + c(s)^2$. Assumption (i) implies $|D(s)| \le C_D\, w(s)$ for some finite C_D. Assumption (ii) implies the drift (Lyapunov) bound

$$\mathbb{E}[w(S') \mid s] = 1 + \mathbb{E}[(\rho c + \varepsilon)^2] \;\le\; \rho^2(1 + c^2) + b \;=:\; \alpha\, w(s) + b,$$

with $\alpha := \rho^2 < 1$ and $b := 1 - \rho^2 + \mu_\varepsilon^2 + \sigma_\varepsilon^2 < \infty$. Note that, in our paper, we assume the shocks to be iid, so Assumption (iii) gives a uniform bound on the default branch: $|-\omega L(s)| \le \omega B^{\max}$. Hence the lemma's conditions hold, so T is a contraction on $(B_w, \|\cdot\|_w)$ with modulus $\beta\alpha < 1$; existence, uniqueness, and convergence follow from Blackwell/Banach (cf. Stachurski 2009, Ch. 12.2, Stokey, Robert E. Lucas, and Prescott 1989 Ch. 3; for an argument with bounded payoffs). □

Corollary 2 (Comparative statics in profits). *Suppose $\bar{\pi}' \ge \bar{\pi}$ implies for all relevant states s: (i) $D_{\bar{\pi}'}(s) \ge D_{\bar{\pi}}(s)$ (larger current profits weakly raise cash flow), and (ii) the nondefault branch weakly improves (e.g., $L^{\max}$ rises with $\bar{\pi}$, shrinking the default region). Then the fixed point is monotone: $V_{\bar{\pi}'}(s) \ge V_{\bar{\pi}}(s)$ for all s.*

Proof. For any V, $(T_{\bar{\pi}'}V)(s) \ge (T_{\bar{\pi}}V)(s)$ pointwise by (i)–(ii). Iterate the two contractions from the same V_0; order is preserved by monotonicity, and limits are the unique fixed points. □

Corollary 3 (Comparative statics in current debt). *Fix $\bar{\pi}$. If two states differ only in L with $L' \ge L$ and (i) $D(s') \le D(s)$, (ii) $-\omega L' \le -\omega L$, and (iii) the nondefault branch is weakly worse at L', then $V(s') \le V(s)$.*

Proof. The three conditions make each branch of TV weakly lower at s' than at s; iterate and pass to the fixed point. □

Corollary 4 (Monotonicity carries to expectations). *Let $\bar{V}(\bar{\pi}, L) := \mathbb{E}[\,V(S) \mid (\bar{\pi}, L, \text{others})\,]$ where the kernel $Q(\cdot \mid s)$ is held fixed. If V is increasing in $\bar{\pi}$ and decreasing in L pointwise (Corollaries 2–3), then $\bar{V}$ is increasing in $\bar{\pi}$ and decreasing in L as well.*

Proof. Monotone functions remain monotone under integration with a fixed probability kernel: if $V_1 \le V_2$ pointwise, then $\mathbb{E}[V_1] \le \mathbb{E}[V_2]$. □

9.8.5 Ex-ante pricing with default.

If hotels have uncertainty over their demand shocks before they set prices, their optimal price (i.e., the price that maximizes lifetime value) overindexes on the states with low shock realizations. The intuition is that hotels lose much more value in default, which is more likely with low demand shocks, than if they set prices too low in states of high demand shocks.

Let $q = q(p, \xi)$ with elasticity $\varepsilon = \partial q / \partial p \cdot p/q < 0$ and $\pi(p, \xi) = (p - mc)\, q(p, \xi)$ so that $\pi_p = q\,(1 + \varepsilon \ell)$, where $\ell = (p - mc)/p$. Let $c \sim D_c$, $B = L_{t-1}(1 + r_{t-1})$, and new-period borrowing be $L(p, \xi, c; B) = B - \pi(p, \xi) + c$. Refinancing feasibility requires $L \leq L_t^{\max}(p, B)$, i.e.

$$g(p, \xi, c; B) \equiv L_t^{\max}(p, B) - B + \pi(p, \xi) - c \;\geq 0.$$

For each c, let $\xi^*(p, c; B)$ solve $g(p, \xi^*, c; B) = 0$: if $\xi \geq \xi^*$ the firm continues, otherwise it defaults.

State payoffs. If the firm continues,

$$A(p, \xi, c) = \pi(p, \xi) - c + \beta\, \mathbb{E}[V_{t+1}(Y)]\,, \qquad Y = (1 + r_t)\, L(p, \xi, c; B).$$

Let $M(p, \xi, c) \equiv \beta\, \mathbb{E}[V_{t+1,Y}(Y)] < 0$ denote the marginal value of next-period liability. If the firm defaults, the payoff is $B^{\text{def}} = -\omega L_{t-1}$ (independent of p).

Ex-ante objective and differentiation. The ex-ante value at price p is

$$W(p) = \iint_{\xi \geq \xi^*(p,c)} A(p, \xi, c)\, f_\xi(\xi) f_c(c)\, d\xi\, dc \;+\; \iint_{\xi < \xi^*(p,c)} B^{\text{def}}\, f_\xi(\xi) f_c(c)\, d\xi\, dc.$$

Using Leibniz' rule with a moving boundary (for each c) and that B^{def} does not depend on p,

$$W_p = \int_{\mathbb{R}} \left[\int_{\xi^*(p,c)}^{\infty} A_p(p, \xi, c)\, f_\xi(\xi)\, d\xi \;+\; \big(B^{\text{def}} - A(p, \xi^*(p,c), c)\big)\, f_\xi\big(\xi^*(p,c)\big)\, \frac{\partial \xi^*(p,c)}{\partial p} \right] f_c(c)\, dc. \tag{21}$$

Compute A_p. Since $L = B - \pi + c$ one has $L_p = -\pi_p$ and thus

$$A_p = \pi_p + M\, \frac{\partial}{\partial p}\big[(1 + r_t) L\big] = \pi_p\, [1 - M(1 + r_t)]. \tag{$\star$}$$

Let the boundary *jump* be $\Delta(p, c) \equiv A(p, \xi^*, c) - B^{\text{def}} \neq 0$. By the implicit function theorem, $\xi^*_p = -g_p/g_\xi$, with

$$g_p = L^{\max}_{t,p}(p,B) + \pi_p(p,\xi^*), \qquad g_\xi = \pi_\xi = (p-mc)\,q_\xi(p,\xi^*).$$

Writing $q_\xi = \kappa\, q$ with $\kappa \equiv \partial \ln q/\partial \xi > 0$ and using $\pi_p = q(1+\varepsilon\ell)$, the boundary contribution in (**??**) equals

$$\int \Delta(p,c)\, f_\xi(\xi^*(p,c))\, \left[\frac{1+\varepsilon\ell}{p\,\ell\,\kappa(p,\xi^*)} + \frac{L^{\max}_{t,p}(p,B)}{p\,\ell\,\kappa(p,\xi^*)\,q(p,\xi^*)}\right] f_c(c)\,dc. \qquad (\star)$$

Collecting terms Define the continuation weight, boundary weight, and the two residual terms:

$$\mathcal{Q} = \iint_{\xi\geq\xi^*} q\left[1 - M(1+r_t)\right] f_\xi f_c\, d\xi\, dc, \qquad \mathcal{B} = \int \frac{\Delta\, f_\xi(\xi^*)}{p\,\ell\,\kappa(\xi^*)}\, f_c(c)\,dc,$$

Collecting terms proportional to $1+\varepsilon\ell$ and the residuals, the first-order condition becomes

$$(1+\varepsilon\ell)\,(\mathcal{Q}+\mathcal{B}) \;+\; +\mathcal{L}_p \;=\; 0. \qquad (\star)$$

Ex-ante markup and "as-if" marginal-cost add-on. Solving for ℓ,

$$\ell^{EA} = -\frac{1}{\varepsilon} - \frac{\mathcal{L}_p}{\varepsilon\,(\mathcal{Q}+\mathcal{B})} = -\frac{1}{\varepsilon}(1+\Phi(B,\bar{\pi})) \qquad (\star)$$

Equivalently, writing the Lerner rule as $1+\varepsilon\big(\ell^{EA} - \Lambda^{EA}/p\big) = 0$, the implied "as-if" marginal-cost wedge share is

$$\frac{\Lambda^{EA}}{p} = \frac{\mathcal{L}_p}{\varepsilon\,(\mathcal{Q}+\mathcal{B})}. \qquad (\star)$$

This markup is higher for firms which are more leveraged, and for firms with lower profit, since the cost shock is large relative to their size.

Intuition for individual terms

- $\mathcal{Q} > 0$ captures the usual pricing force in continuation states, scaled down by $1-M(1+r_t) \in (0,1]$ because selling more today raises tomorrow's debt (and $M < 0$).

- $\mathcal{B} \geq 0$ is a boundary mass that increases the *effective* weight on the Lerner kernel.

- and $\mathcal{L}_p$ adds a wedge when the debt limit relaxes with price ($L^{\max}_{t,p} \geq 0$).

Thus $\Lambda^{EA}/p > 0$ i.e. firms behave as if marginal cost were higher.

Shadow marginal cost. Let $\varepsilon < 0$ be the price elasticity at the chosen price p. Without credit risk, inverting the demand system recovers the "marginal cost" $\widehat{mc} = p\big(1 + \frac{1}{\varepsilon}\big)$.

However, with financing/default risk, the ex-ante pricing condition can be written as a Lerner rule with an added wedge summarized by $\Delta > 0$, yielding

$$mc^{\text{shadow}} \; = \; p\Big(1 + \frac{1+\Delta}{\varepsilon}\Big)$$

i.e., the inferred marginal cost equals the standard $p\big(1 + \frac{1}{\varepsilon}\big)$ *plus* an additional term $\frac{\Delta}{\varepsilon}\, p$.

The wedge Δ rises with leverage (B), with the dispersion of unpriced demand shocks (D_{ξ}), with the dispersion of cost shocks (D_c), and with the interest rate and debt limit.

Ex-ante markup and "as-if" marginal-cost add-on. Solving for ℓ,

$$\ell^{EA} = -\frac{1}{\varepsilon} - \frac{\mathcal{R} + \mathcal{L}_p}{\varepsilon\,(\mathcal{Q} + \mathcal{B})}. \tag{$\star$}$$

Equivalently, writing the Lerner rule as $1 + \varepsilon\big(\ell^{EA} - \Lambda^{EA}/p\big) \; = \; 0$, the implied "as-if" marginal-cost wedge share is

$$\frac{\Lambda^{EA}}{p} = \frac{\mathcal{R} + \mathcal{L}_p}{\varepsilon\,(\mathcal{Q} + \mathcal{B})}. \tag{$\star$}$$

9.9 Model Equilibrium

Market averages and credit profiles. Let Ψ denote the market-average state summarizing the aggregate conditions used to form expectations (e.g., the number of active hotels, entry and exit probabilities, and average profitability). We shall suppress the dependence on time for brevity, but note that these conditions need to hold for all periods t.

Each hotel j has a *credit profile*

$$c_j \; = \; (\bar{\pi}_j, \mathcal{H}_j),$$

where $\bar{\pi}_j$ is its expected mean operating profit and $\mathcal{H}_j$ summarizes its operating history (the statistic used to update beliefs about quality).

Approval rule. Competitive lenders share a common approval rule

$$\mathcal{L} : \; \mathcal{C} \to \mathbb{R}_+, \qquad L_j^{\max} \; = \; \mathcal{L}(c_j),$$

which assigns to each profile $c \in \mathcal{C}$ the largest loan size consistent with zero expected profit.

Lenders: Zero-Profit and the Credit Forward Map For a proposed loan L to a borrower with profile c, and a belief about the next-period approval rule $\mathcal{L}'$ (primes denote next-period objects), the lender's expected profit at the zero-profit rate r is

$$\mathbb{E}\Big[\pi^{\text{lender}}(r, L, c \mid \mathcal{L}')\Big] \;=\; (1-p)\,(1+r)L \;+\; p\,\omega L \;-\; (1+m)L \;=\; 0, \tag{22}$$

where m is the cost of funds, $\omega \in (0,1]$ is the recovery rate upon default, and

$$p \;=\; \Pr\big\{\pi' \;\leq\; L(1+r) \;-\; \mathcal{L}'(c')\big\} \tag{23}$$

is the default probability, which depends on the borrower's refinancing ability next period under $\mathcal{L}'$.

A loan L is feasible if there exists an r satisfying (22).

Backward induction operator. Define the operator

$$T_{\text{credit}} : \; (\mathcal{C} \to \mathbb{R}_+) \;\longrightarrow\; (\mathcal{C} \to \mathbb{R}_+),$$

which takes the believed future rule $\mathcal{L}'$ and returns the *current* rule given lender beliefs:

$$T_{\text{credit}}[\mathcal{L}'](c) \;:=\; \sup\Big\{L \geq 0 : \; \exists\, r \text{ s.t. } \mathbb{E}\Big[\Pi^{\text{lender}}(r, L, c; \mathcal{L}')\Big] = 0\Big\}. \tag{24}$$

In a stationary equilibrium, beliefs are correct and the rule reproduces itself:

$$\mathcal{L} \;=\; T_{\text{credit}}[\mathcal{L}]\,. \tag{25}$$

Prices, Profits, Rollover, Default/Exit, and Entry

(1) Prices (accommodation market). Given $\bar{\Psi}$, active hotels play Nash–Bertrand in differentiated products (nested-logit demand). For each j,

$$p_j^* \;=\; \arg\max_{p_j} \; s_j(\mathbf{p}; \Psi)\,\big(p_j - c_j^{\text{mc}}\big), \tag{26}$$

yielding the equilibrium price vector $\mathbf{p}^*$, shares s_j^*, revenue $R_j = M\, p_j^* s_j^*$ (market size M), and operating profit

$$\pi_j \;=\; \bar{\pi}_j \;+\; \pi_j^{\xi} \;-\; \mathcal{C}_j \;+\; \pi_j^{\phi}, \tag{27}$$

with the distribution of $(\pi_j^{\xi}, \mathcal{C}_j, \pi_j^{\phi})$ implied by $\bar{\Psi}$ and the demand/cost shock processes.

(2)Belief consistency (Bayesian updating). Lenders share a common prior μ_0 over the hidden type ϕ_j (full support). Here we specify a time subscript for clarity.

For each firm j and date t, given history $\mathcal{H}_{jt}$ and the model-implied likelihood $f(\pi_{jt} \mid \phi_j, \mathcal{H}_{jt})$, lenders' equilibrium beliefs are the Bayes posterior:

$$\mu_{j,t}(\phi \mid \mathcal{H}_{jt}) \;=\; \frac{f(\pi_{jt-1} \mid \phi, \mathcal{H}_{jt-1})\, \mu_{jt-1}(\phi \mid \mathcal{H}_{jt-1})}{\int f(\pi_{jt-1} \mid \tilde{\phi}, \mathcal{H}_{jt-1})\, \mu_{jt-1}(\tilde{\phi} \mid \mathcal{H}_{jt-1})\, d\tilde{\phi}}, \qquad t \geq 1,$$

with $\mu_{j,0}(\phi \mid \mathcal{H}_{j,0}) = \mu_0(\phi)$.

(3) Debt accounting, rollover, dividends. Let last period's approved loan be L_j at rate r_j; the current obligation is

$$B_j \;=\; (1 + r_j)\, L_j. \tag{28}$$

After π_j realizes:

$$\text{if } \pi_j \geq B_j : \quad D_j = \pi_j - B_j, \;\text{ no new credit requested;} \qquad \text{if } \pi_j < B_j : \quad L_j^{\text{req}} = B_j - \pi_j.$$

If approved, the next obligation is $B_j' \;=\; (1 + r_j')\, L_j^{\text{req}}$, where r_j' is the zero-profit rate associated with (L_j^{req}, c_j').

(4) Default and exit. Approval uses the rule $\mathcal{L}$. Default occurs iff the required rollover exceeds the approved limit:

$$\mathbf{1}\{\text{default}_j\} \;=\; \mathbf{1}\Big\{\, L_j^{\text{req}} \;>\; \mathcal{L}(c_j') \,\Big\}. \tag{29}$$

Upon default, the firm exits and the lender recovers ωL_j^{req}. Hence default and exit coincide.

(5) Entry. Potential entrants draw (F, c_0, η) from $\mathcal{D}_{\text{entry}}$, finance $(1 - s)F$ with debt, and post equity sF. Let $V(c, B; \bar{\Psi}, \mathcal{L})$ be the incumbent value. An entrant enters iff

$$\mathbb{E}\big[V(c_0, B_0; \bar{\Psi}, \mathcal{L})\big] \;-\; sF \;\geq\; 0, \qquad B_0 \;=\; (1 + r_0(\mathcal{L}))\,(1 - s)F. \tag{30}$$

Value recursion With transitions $c \to c'$ (Bayesian updating from $\mathcal{H}$), and $B' = (1 + r')\max\{0, B - \pi\}$,

$$\begin{aligned} V(c, B; \bar{\Psi}, \mathcal{L}) \;=\; \mathbb{E}\big[\, & \mathbf{1}\{\pi \geq B\} \left(\pi - B + \beta V(c', B'; \bar{\Psi}, \mathcal{L})\right) \\ & +\; \mathbf{1}\{\pi < B\}\, \mathbf{1}\{L^{\text{req}} \leq \mathcal{L}(c')\}\, \beta V(c', B'; \bar{\Psi}, \mathcal{L}) \\ & -\; \mathbf{1}\{L^{\text{req}} > \mathcal{L}(c')\}\, \omega\, L^{\text{req}} \big]. \end{aligned} \tag{31}$$

Market-Structure Map Given $(\Psi, \mathcal{L})$, within-period outcomes aggregate to next period's market averages via the operator

$$\mathcal{G} :\; (\Psi, \mathcal{L}) \longmapsto \Psi',$$

defined by

$$\Psi' \;=\; \mathcal{G}(\Psi; \mathcal{L}) \;=\; \Phi\Big(e(\Psi; \mathcal{L}),\; \delta(\Psi; \mathcal{L}),\; \text{pricing outcomes under } \Psi\Big), \tag{32}$$

where $\Phi(\cdot)$ summarizes how entry, exit/default, and price competition update the market-average statistics.

Stationary Fixed Points (Two Maps) A *stationary Markov (oblivious) rational-expectations equilibrium* is a pair $(\Psi^*, \mathcal{L}^*)$ such that

$$\mathcal{L}^* \;=\; T_{\text{credit}}[\mathcal{L}^*], \qquad \Psi^* \;=\; \mathcal{G}(\Psi^*; \mathcal{L}^*). \tag{33}$$

All within-period outcomes—equilibrium prices $\mathbf{p}^*$, operating profits π, dividends D, the default/exit rate $\delta(\Psi^*; \mathcal{L}^*)$, and the entry rate $e(\Psi^*; \mathcal{L}^*)$—are those induced by $(\Psi^*, \mathcal{L}^*)$.

9.10 Estimation Details

9.10.1 Proofs

Instruments across Time We will show the conditions under which estimation with lead prices, instrumented with instrument z_{t+2} gives a valid unbiased estimate of price sensitivity β.

The structural equation is

$$\tilde{y}_t \;=\; \beta\,\tilde{p}_t + \tilde{u}_t,$$

but empirically we only observe $\tilde{p}_{t+2}$ with

$$\tilde{p}_{t+2} \;=\; \tilde{p}_t + \tilde{\iota}_{t+2}.$$

We use $\tilde{z}_{t+2}$ as an instrument.

Assumption 1 (Relevance). $\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{p}_t) \neq 0$.

Assumption 2 (Exclusion). $\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{u}_t) = 0$.

Assumption 3 (Evolution orthogonality). $\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{\iota}_{t+2}) = 0$.

Proposition 4. *Under (1), (2), (3), IV 2SLS estimate using z_{t+2} as an instrument for p_{t+2} is unbiased.*

Proof. If we estimate

$$\tilde{y}_t = \beta\,\tilde{p}_{t+2} + \tilde{\varepsilon}_t, \quad \text{where} \quad \tilde{\varepsilon}_t := \tilde{u}_t - \beta\,\tilde{\iota}_{t+2},$$

by IV using $\tilde{z}_{t+2}$, the population estimand is

$$\beta^{\mathrm{IV}} = \frac{\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{y}_t)}{\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{p}_{t+2})} = \frac{\mathrm{Cov}(\tilde{z}_{t+2},\, \beta\tilde{p}_{t+2} - \beta\tilde{\iota}_{t+2} + \tilde{u}_t)}{\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{p}_{t+2})} = \beta + \frac{\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{u}_t) - \beta\,\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{\iota}_{t+2})}{\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{p}_{t+2})}.$$

Under Assumptions 2 and 3, the numerator of the bias term is zero, hence

$$\beta^{\mathrm{IV}} = \beta \quad \text{provided} \quad \mathrm{Cov}(\tilde{z}_{t+2}, \tilde{p}_{t+2}) \neq 0.$$

□

If 3 fails, the asymptotic bias is

$$\beta^{\mathrm{IV}} - \beta = -\,\beta\,\frac{\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{\iota}_{t+2})}{\mathrm{Cov}(\tilde{z}_{t+2}, \tilde{p}_{t+2})},$$

whose sign is $\operatorname{sign}\big(-\beta\cdot\operatorname{Cov}(\tilde z_{t+2},\tilde\iota_{t+2})\big)$.

Remark (ideal but infeasible baseline). If instead one could estimate the correct equation with $\tilde p_t$ on the right-hand side and instrument $\tilde p_t$ by $\tilde z_{t+2}$, then

$$\beta^{\mathrm{IV}}=\frac{\operatorname{Cov}(\tilde z_{t+2},\,\beta\tilde p_t+\tilde u_t)}{\operatorname{Cov}(\tilde z_{t+2},\tilde p_t)}=\beta+\frac{\operatorname{Cov}(\tilde z_{t+2},\tilde u_t)}{\operatorname{Cov}(\tilde z_{t+2},\tilde p_t)}=\beta\quad\text{under Assumptions 1–2},$$

and no additional restriction on $\operatorname{Cov}(\tilde z_{t+2},\tilde\iota_{t+2})$ is needed beyond ensuring a nonzero first stage via $\operatorname{Cov}(\tilde z_{t+2},\tilde p_t)\neq 0$.

Profit Cutoff for Entry

Cutoff formulation. Fix a (possibly π–dependent) fixed cost $L(\pi)\geq 0$ and define

$$G(\pi)\;:=\;V\big(\pi,\,L(\pi)\big)\;-\;L(\pi),$$

where $V(\pi,L)$ is the unique bounded value function from the dynamic problem (cf. earlier results), and recall that V is (weakly) increasing in the current-period profit index π and (weakly) decreasing in L.

Proposition 5 (Cutoff equivalence)**.** *Suppose G is (strictly) increasing in π and there exist $\underline{\pi}<\overline{\pi}$ such that $G(\underline{\pi})<0<G(\overline{\pi})$. Then there exists a (unique) cutoff $\pi^*\in(\underline{\pi},\overline{\pi})$ with*

$$G(\pi^*)=0,\qquad\textit{and}\qquad G(\pi)\geq 0\iff\pi\geq\pi^*.$$

Equivalently,

$$V\big(\pi,L(\pi)\big)\;-\;L(\pi)\;\geq\;0\quad\iff\quad\pi\;\geq\;\pi^*.$$

Proof. Since G is strictly increasing and changes sign between $\underline{\pi}$ and $\overline{\pi}$, the Intermediate Value Theorem yields a unique root π^* with $G(\pi^*)=0$. Monotonicity of G then implies $G(\pi)\geq 0$ iff $\pi\geq\pi^*$. □

Corollary 5 (Constant-L case)**.** *If $L(\pi)\equiv\bar L$ (constant), then $G(\pi)=V(\pi,\bar L)-\bar L$ is (weakly) increasing in π because V is (weakly) increasing in π. Whenever G crosses zero, a (unique) cutoff π^* exists and $V(\pi,\bar L)-\bar L\geq 0\;\Leftrightarrow\;\pi\geq\pi^*$.*

Remark 1 (Proportional leverage)**.** *For proportional borrowing $L(\pi)=\lambda\,\pi$ with $\lambda\geq 0$, we adopt the cutoff formulation by* assuming *(and then verifying in computations across many λ) that $G(\pi)=V(\pi,\lambda\pi)-\lambda\pi$ is increasing in π and crosses zero once. Under these conditions, Proposition 5 applies, delivering a cutoff rule $V(\pi,\lambda\pi)-\lambda\pi\geq 0\;\Leftrightarrow\;\pi\geq\pi^*$.*

2-Step Estimation of Fixed Cost Using the cutoff lemma above, we can estimate the fixed cost L in two steps.

Step 1 (Cutoff in profits). The entry decision follows a threshold rule:

$$g_L(\pi + \eta) > 0 \iff \pi + \eta > \tau(L).$$

Equivalently,

$$V(\pi + \eta, (1-c)L) - cL > 0 \iff \pi + \eta > \pi^\star,$$

so Step 1 identifies the cutoff $\pi^\star = \tau(L)$ from the observed entry/no–entry outcomes.

Step 2 (Criterion for L). Define the indicator $\mathbf{1}\{F(\pi + \eta, L) > 0\}$ for entry. By monotonicity of F in π, there exists a unique cutoff $c(L)$ such that

$$F(c(L), L) = 0.$$

Hence

$$\mathbf{1}\{F(\pi + \eta, L) > 0\} \;=\; \mathbf{1}\{\pi + \eta > c(L)\} \;=\; \mathbf{1}\{G + \eta > 0\}, \quad G := \pi - c(L).$$

The population likelihood is therefore

$$l(L) \;=\; \mathbb{E}[\mathbf{1}\{G + \eta > 0\}] \;=\; \mathbb{E}[\Pr(\eta > -G \mid G)]\,.$$

If $\eta \sim \mathcal{N}(0, \sigma^2)$ and is independent of π, then

$$\Pr(\eta > -G \mid G) = \Phi\Big(\tfrac{G}{\sigma}\Big)\,,$$

so

$$l(L) = \mathbb{E}\Big[\Phi\Big(\tfrac{\pi - c(L)}{\sigma}\Big)\Big]\,.$$

Reduction. By construction $c(L)$ is defined by $F(c(L), L) = 0$. Once Step 1 yields the cutoff $\pi^\star$, the admissible pair $(\pi^\star, L)$ must satisfy

$$F(\pi^\star, L) = 0.$$

If F is strictly decreasing in L, then $L \mapsto c(L)$ is strictly monotone and hence invertible, so there exists at most one L^* with

$$F(\pi^\star, L^*) = 0.$$

Thus, the Step 2 estimator reduces to solving this single equation for L^*.

Remark. Normality of η is not required for the cutoff–inversion logic. The normal assumption is only used to express $l(L)$ in closed form; any distribution of η with strictly increasing cdf suffices for identification.

Posterior Type from Signal Formally, lenders update beliefs about firm type ϕ using Bayesian inference on profit histories. With normal shocks to demand $\xi \sim \mathcal{N}(0, \sigma_\xi^2)$, Bayesian updating delivers closed-form posteriors.

Posterior. Given$\mathcal{H} = \{\pi_i\}_{i=1}^T$, $(\bar{\pi} + \phi_{post}) \sim \mathcal{N}(\mu_T,\ \sigma_T^2)$,

$$\sigma_T^{-2} \ = \ \sigma_0^{-2} \ + \ \sum_{i=1}^{T} \frac{1}{\tau_i^2}, \quad \mu_T \ = \ \sigma_T^2 \left(\frac{\mu_0}{\sigma_0^2} \ + \ \sum_{i=1}^{T} \frac{\pi_i}{\tau_i^2} \right).$$

where

$$\mu_T \ = \ \frac{\sum_{i=1}^T w_i R_i}{\sum_{i=1}^T w_i}, \qquad \sigma_T^2 \ = \ \frac{1}{\sum_{i=1}^T w_i}, \qquad w_i \ := \ \frac{1}{(\sigma_\phi R_i)^2 + \sigma_\xi^2}.$$

As firms survive longer, type uncertainty shrinks, and inferred ϕ converges toward the true type. Greater noise in shocks slows this convergence.

Note that, since the variance scales with profit, we need an approximation to obtain posteriors as approximately normal

Lemma 3 (Local Gaussian posterior with proportional noise)**.** *Let a scalar type μ (mean profit) have an (improper) flat prior. Suppose we observe R with heteroskedastic Gaussian noise,*

$$R \mid \mu \sim \mathcal{N}\big(\mu,\ (\sigma_\phi\, \mu)^2\big), \qquad \sigma_\phi > 0.$$

Then for realizations $R > 0$ and in a local neighborhood where $\mu \approx R$, the posterior admits the Laplace (quadratic) approximation

$$\mu \mid R \ \approx \ \mathcal{N}\big(R,\ (\sigma_\phi\, R)^2\big).$$

Proof. With a flat prior, the posterior is proportional to the likelihood

$$\ell(\mu; R) \;\propto\; \frac{1}{\sigma_\phi\,\mu}\,\exp\!\left(-\frac{(R-\mu)^2}{2(\sigma_\phi\,\mu)^2}\right), \qquad \mu > 0.$$

Let $g(\mu) = \log \ell(\mu; R)$. A second–order Taylor expansion of g about the mode $\hat{\mu}$ (which is $\hat{\mu} = R + o(1)$ when $R > 0$ and σ_ϕ is small) yields

$$g(\mu) \;\approx\; g(\hat{\mu}) \;-\; \frac{(\mu-\hat{\mu})^2}{2\,(\sigma_\phi\,R)^2},$$

where we (i) plug in $\mu \approx R$ inside the variance term, and (ii) note that the curvature from the prefactor $-\log \mu$ is $O(1/R^2)$ and thus dominated by the quadratic term when R is not too small or σ_ϕ is modest. Exponentiating gives the stated normal approximation. $\square$

(i) The result can be read as: with a flat prior, the posterior is (locally) the likelihood, and if the noise s.d. scales proportionally with the level, one can freeze the scale at the realized R to obtain a Gaussian $\mathcal{N}(R, (\sigma_\phi R)^2)$.

Lemma 4. *The posterior update on a firm's type with history $\mathcal{H}$ of profits, with gaussian types and gaussian demand shocks is approximately gaussian.*

Proof. This follows from the above lemma, where the posterior from each location experiment is approximately normal. Then given these signals:

$$R_i \mid \mu \;\approx\; \mathcal{N}(\mu,\, \tau_i^2), \qquad \tau_i^2 \;:=\; (\sigma_\phi R_i)^2 + \sigma_\xi^2, \quad i = 1, \ldots, T.$$

we have the following posterior:

Posterior. $\mu \mid \{R_i\}_{i=1}^T \;\approx\; \mathcal{N}(\mu_T,\, \sigma_T^2), \quad \sigma_T^{-2} \;=\; \sigma_0^{-2} \;+\; \sum_{i=1}^T \frac{1}{\tau_i^2}, \quad \mu_T \;=\; \sigma_T^2\left(\frac{\mu_0}{\sigma_0^2} \;+\; \sum_{i=1}^T \frac{R_i}{\tau_i^2}\right).$

where

$$\mu_T \;=\; \frac{\sum_{i=1}^T w_i R_i}{\sum_{i=1}^T w_i}, \qquad \sigma_T^2 \;=\; \frac{1}{\sum_{i=1}^T w_i}, \qquad w_i \;:=\; \frac{1}{(\sigma_\phi R_i)^2 + \sigma_\xi^2}.$$

$\square$

Local approximation from nested logit to multiplicative profits. In nested logit, product j in nest g has mean utility δ_{jt} and (idiosyncratic) additive shock ξ_{jt} entering $\delta_{jt} + \xi_{jt}$.

The market share factors as

$$s_{jt}=s_{j|g,t}\,s_{g,t},\quad s_{j|g,t}=\frac{\exp\{\delta_{jt}/(1-\sigma)\}}{\sum_{k\in g}\exp\{\delta_{kt}/(1-\sigma)\}},\quad s_{g,t}=\frac{\left(\sum_{k\in g}\exp\{\delta_{kt}/(1-\sigma)\}\right)^{1-\sigma}}{1+\sum_h\left(\sum_{k\in h}\exp\{\delta_{kt}/(1-\sigma)\}\right)^{1-\sigma}}.$$

Perturb $\delta_{jt}\mapsto\delta_{jt}+\xi_{jt}$. Using $s_{jt}=s_{j|g,t}s_{g,t}$ and differentiating $\log s_{jt}$ with respect to δ_{jt} gives

$$\frac{\partial\log s_{jt}}{\partial\delta_{jt}}=\underbrace{\frac{1}{1-\sigma}\left(1-s_{j|g,t}\right)}_{\text{within-nest effect}}+\underbrace{\left(1-s_{g,t}\right)s_{j|g,t}}_{\text{across-nest effect}}\;\equiv\;\kappa_{jt}.$$

Thus, for a shock ξ_{jt},

$$\log s_{jt}(\xi_{jt})\;\approx\;\log\bar{s}_{jt}+\kappa_{jt}\,\xi_{jt}\quad\Rightarrow\quad s_{jt}(\xi_{jt})\;\approx\;\bar{s}_{jt}\,e^{\kappa_{jt}\xi_{jt}}.$$

With $R_{jt}=p_{jt}M_ts_{jt}$ and $\pi_{jt}\equiv m_{jt}R_{jt}$, holding p_{jt} and M_t fixed (or absorbed by time effects) and taking markups m_{jt} locally constant,

$$\log\pi_{jt}(\xi_{jt})\;\approx\;\log\bar{\pi}_j+\kappa_{jt}\,\xi_{jt}\quad\Rightarrow\quad\pi_{jt}(\xi_{jt})\;\approx\;\bar{\pi}_j\,e^{\kappa_{jt}\xi_{jt}}.$$

Hence the first-order approximation is $\pi_{jt}\approx\bar{\pi}_j(1+\kappa_{jt}\xi_{jt})$, where the multiplier

$$\kappa_{jt}=\frac{1}{1-\sigma}(1-s_{j|g,t})+(1-s_{g,t})s_{j|g,t}$$

is the semi-elasticity of $\log s_{jt}$ to an additive utility shock and depends on the nested-logit substitution structure (shares and σ).

If $\xi_{jt}\sim\mathcal{N}(0,\sigma_\xi^2)$ and κ_{jt} varies slowly across j,t (or is well-approximated by a common κ), the implied multiplicative form yields a constant (or approximately constant) coefficient of variation:

$$\mathrm{CV}(\pi_{jt})=\sqrt{e^{\kappa_{jt}^2\sigma_\xi^2}-1}\;\approx\;\text{constant across firms.}$$

Local Linear Approximation from Demand Shocks We estimate our model, assuming that the additive profit terms π_ξ and π_ϕ scale linearly with mean profit $\bar{\pi}$. The following lemma can be used as a basis for internal consistency, so that the following linear assumptions are consistent with our demand model.

Lemma 5 (Linearized revenue response)**.** *For a perturbation ε to product i's mean utility, the revenue admits the first–order approximation*

$$R_i(\delta_i+\varepsilon,\delta_{-i})\;\approx\;R_i(\delta_i,\delta_{-i})\;\left[1+\lambda_{it}\,\varepsilon\right],$$

where

$$\lambda_{it} = \frac{\partial \ln s_{it}}{\partial V_{it}} = \frac{1 - \sigma s_{i|g,t} - (1-\sigma) s_{i|g,t} S_{g(t),t}}{1-\sigma}.$$

For small $s_{i|g,t}$ and s_{it},

$$\lambda_{it} = \frac{1}{1-\sigma} - \frac{\sigma}{1-\sigma} s_{i|g,t} - s_{it} + o\Big(s_{i|g,t}, s_{it}\Big).$$

Setting $\sigma = 0.65$ yields

$$\lambda_{it} \approx \frac{1}{0.35} = 2.857 \quad \text{with first-order correction} \quad -1.857\, s_{i|g,t} - s_{it}.$$

In particular, if $s_{i|g,t} \leq 0.05$ and $s_{it} \ll s_{i|g,t}$, then $\lambda_{it} \in [2.76, 2.86]$.

With $s_{i|g,t} \approx 1/n_g$,

$$\lambda_{it} \;\approx\; \frac{1}{1-\sigma} \;-\; \frac{\sigma}{(1-\sigma)\, n_g} \;-\; \frac{S_{g(t),t}}{n_g}.$$

This approximation represents our markets well.

Note that we impose this linearity during estimation, and while running counterfactuals, we draw profit perturbations using these estimated linear models. This means that the empirical setup is fully self-consistent across our estimation and results. Through above, we argue that this assumption is also approximately consistent with our demand model.

9.10.2 Inverting Revenues to Obtain Prices

Let R_j denote firm j's revenue and p_j represent its price, which is the object of interest. We will suppress dependence on time.

For 2010–2019, prices are solves from from revenues by solving:

$$s_j(\mathbf{p}) = \frac{R_j/p_j}{\sum_k R_k/p_k}, \qquad p_j - mc_j = -\frac{s_j(\mathbf{p})}{\partial s_j(\mathbf{p})/\partial p_j}, \tag{34}$$

which defines J equations in J unknown prices.

We solve the system via dampened Newton's method, where a candidate vector of prices is used to obtain shares as given above. Using those shares, we obtain an implied vector of prices from the markup equation. Equilibrium prices are such that the two vectors are (numerically).

We solve for the zero of this fixed point using Dampened Netwon's method, where the jacobian is calculated approximately using numerical derivatives. We use a dampening parameter of 0.6. The procedure achieves convergence in over 98% of market-years within 20 iterations.

9.10.3 Model Solution Details

Stage 1: Solving for credit terms We discretize the state space along four dimensions: $(\sigma_\phi, \sigma_c, \pi, t)$, where t indexes firm age from entry through ten periods of operation. For the uncertainty parameters, we use a 20×20 grid: σ_c ranges from 0 to the 25th percentile of mean profit (capturing the scale of operational cost shocks), while σ_ϕ ranges from 0 to 0.8 standard deviations of mean profit (capturing dispersion in permanent types). For the profit dimension, we use 40 points spanning the empirical profit distribution, with 20 points concentrated in the bottom quartile to capture small firms more precisely. For the type dimension, we approximate the continuous distribution using three-point Gaussian quadrature, so each profit level has three possible associated types. With all cost parameterizations being normal shocks, the implied posterior distribution over types remains approximately normal (see Appendix (9.10.1).

For each combination of profit, type, and age (across the $\sigma_\phi - \sigma_c$ grid), we compute equilibrium debt limits and interest rate schedules. Both credit terms vary with firm age as lenders learn about types through observed profit histories. Numerical solution details are provided in Appendix (9.10.3)

Now we provide some computational details of the solution for each profit.

Debt Limits and Interest Rates The model for debt limits and interest rates is solved separately for firms with different π. We start by a grid of firms, which is denser for smaller firms. Then, model solution proceeds backwards in the following steps:

1 First, we solve a single fixed point problem where $L^{''} = L^{'}$. This is analogous to a no-ponzi condition, in the infinite future, firm debt limits stay bounded as all uncertainty is revealed, and there is no difference between lenders' information set

across periods. [52]

$$\mathbb{E}[\pi^{\text{bank}}(r^*, L', L^{'}, \pi')]\Big|_{L'=L_j'^{\max}} = mL',$$

2 Then, assuming debt limit for the t, we solve for the debt limit with period $t-1$, assuming the lenders' distribution of uncertainty around firm is 9.10.1, with lenders' weighted mean being this profit.

3 Proceed backwards, until we have period 0 debt limits.

Stage 2: Debt transitions and exit probabilities : Given credit terms, we construct the debt transition kernel and compute exit probabilities. The kernel describes how the firm's debt evolves stochastically over time: from any current debt level, profit realizations determine whether the firm repays, rolls over debt, or defaults. By tracking these transitions, we can compute the probability that a firm starting with given initial debt will exit in any future period.

The key object is $K_{jt}(L_{jt+1} \mid L_{jt})$, which describes how current loan demand L_{jt} evolves to next-period loan demand L_{jt+1}:

$$K_{jt}(L_{jt+1} \mid L_{jt}) = \begin{cases} 1 - F_{\pi_{jt}}((1+r_{jt})L_{jt}), & \text{if } L_{jt+1} = 0, \\ f_{\pi_{jt}}((1+r_{jt})L_{jt} - L_{jt+1}), & \text{if } L_{jt+1} > 0 \end{cases} \tag{35}$$

where $f_{\pi_{jt}}$ is the density of realized profit. If profit $\pi_{jt} \geq (1+r_{jt})L_{jt}$, the firm repays in full and demands no new loan; otherwise, it rolls over the shortfall $L_{jt+1} = (1+r_{jt})L_{jt} - \pi_{jt}$.

Now, we can recursively apply the transition kernel to compute the distribution of debt in future period t from an initial debt L_{t_0} in period t_0. Let $\Gamma_{jt}(L \mid L_{t_0})$ denote the probability that the firm has debt level L in period t, conditional on starting with L_{t_0} and surviving to t. This evolves recusively according to:

$$\Gamma_{j,t+1}(L' \mid L_{t_0}) = \int K_{jt}(L' \mid L) \cdot \Gamma_{jt}(L \mid L_{t_0}) dL. \tag{36}$$

Finally, we can calculate the period-t exit probability is the probability that debt exceeds the credit limit:

$$p_{jt}(L_{t_0}) = \int \Gamma_{jt}(L \mid L_{t_0}) \cdot \mathbf{1}\{L > \bar{L}_{jt}\}, dL, \tag{37}$$

[52] This follows from the posterior we derived above. Numerically, this is solved as a zero for the following equation, where r^* solves for the firm first order condition, also solved using root-finder.

where $\bar{L}_{jt}$ is the period t credit limit. This captures how firm characteristics (age, profit history, current debt) interact with structural uncertainty to determine exit risk.

Stage 3: Firm value functions. The lender's problem provides two inputs for computing firm values: (i) interest rate schedules that map loan demand into obligations $B_{jt+1} = (1 + r_{jt})L_{jt}$, and (ii) the transition kernel from equation (35). We solve for values in two sub-steps.

First, holding market structure fixed, we compute the value function on a finite grid of (π, B) states using a contraction mapping (Appendix 9.8.4). A key computational insight exploits linearity: instead of solving at each profit realization, we solve for the *ex-ante* expected value function, which reduces to a system of linear equations solvable in closed form (cf. Hotz and Miller (1993b)). This is described below.

Linear Solution to (Expected) Value Function The value of a firm with mean period profit π_j, outstanding debt B_{jt}, and realized cost shock $\mathcal{C}_{jt}$ is

$$V(\pi_j, B_{jt}, \mathcal{C}_{jt}) \;=\; \max\{0,\ \pi_j - B_{jt} - \mathcal{C}_{jt}\} \;+\; \beta\, \mathbb{E}\big[\, V(\pi_j, B_{j,t+1}, \mathcal{C}_{j,t+1}) \,\big],$$

where next period debt is

$$B_{j,t+1} = (1 + r)\ \max\{0,\ B_{jt} + \mathcal{C}_{jt} - \pi_j\}.$$

Taking expectations yields a fixed point in the expected value:

$$\mathbb{E}V(\pi_j, B_{jt}, \mathcal{C}_{jt}) = \int \Big(\max\{0,\ \pi_j - B_{jt} - c\} + \beta\, \mathbb{E}V(\pi_j, B_{j,t+1}(B_{jt}, c), \mathcal{C}_{j,t+1}) \Big)\, f_{\mathcal{C}}(c)\, dc.$$

To compute this object, we discretize debt and shocks. Suppress time and write $EV_j(B) := \mathbb{E}V(\pi_j, B, \mathcal{C})$. Let $\{B_k\}_{k=1}^{K}$ be a grid for debt and let $\{(c_m, w_m)\}_{m=1}^{M}$ be a discrete approximation to $\mathcal{C}$ with weights $\sum_m w_m = 1$. Define the debt transition mapping

$$\mathcal{T}(B, c) \;:=\; (1 + r)\ \max\{0,\ B + c - \pi_j\}.$$

Then the finite-grid approximation to the expected value at B_k is

$$EV_j(B_k) = \sum_{m=1}^{M} w_m \Big[\max\{0,\ \pi_j - B_k - c_m\} + \beta\, EV_j\big(\mathcal{T}(B_k, c_m)\big) \Big].$$

Let $\Lambda(\mathbf{B}) \in \mathbb{R}^K$ collect the current-payoff terms with kth element

$$\Lambda_k \; := \; \sum_{m=1}^{M} w_m \; \max\{0, \; \pi_j - B_k - c_m\},$$

and let $W \in \mathbb{R}^{K \times K}$ be the (row-stochastic) transition matrix induced by $\mathcal{T}$ and the shock weights:

$$W_{k\ell} \; := \; \sum_{m=1}^{M} w_m \; \mathbb{1}\big\{\mathcal{T}(B_k, c_m) \in \mathcal{B}_\ell\big\},$$

where $\{\mathcal{B}_\ell\}$ partitions the debt grid (with interpolation if desired). Stacking $EV_j(B_k)$ into the vector $EV(\mathbf{B}) \in \mathbb{R}^K$, the fixed point is linear:

$$EV(\mathbf{B}) \; = \; \Lambda(\mathbf{B}) \; + \; \beta \, W \, EV(\mathbf{B}).$$

Hence the expected value admits the closed form

$$EV(\mathbf{B}) \; = \; \big(I - \beta W\big)^{-1} \, \Lambda(\mathbf{B}).$$

Intuitively, the value is the discounted sum of current payoffs Λ propagated through the debt-transition matrix W. This expected-value formulation delivers an analytical solution for the finite-grid approximation, in the spirit of expected value iteration (cf.Hotz and Miller 1993b).

Cost MLE via Differential Evolution and BFGS We maximize this using global search via differential evolution to avoid local optima, followed by local refinement with BFGS for precise convergence. Standard errors come from a parametric bootstrap with 100 replications, where we re-draw π^* and re-solve the model in each iteration.

9.11 Non Parametric Identification

We show some results under which the underlying distribution of costs $\mathcal{C}$ is non-parametrically identified.

The core idea behind these identification results is simple. Default probabilities and debt transitions are threshold rules: once you know the debt level and revenues, whether the firm survives depends only on shocks crossing a cutoff. By varying debt L or profits π, we shift that cutoff and therefore trace out parts of the underlying shock distribution $\mathcal{C}$. As profits π grow arbitrarily large, the debt limit $L^{\max}(\pi)$ expands, so the observable support of $\mathcal{C}$ also expands until the full distribution is revealed. With symmetry, the negative side is pinned down as well. In short, debt dynamics act like repeated "location experiments" on $\mathcal{C}$, and observing how exit probabilities respond to different (L, π) pairs is enough to non-parametrically recover the shock distribution.

Theorem 1. If $L^{\max}(\pi) \to \infty$ as $\pi \to \infty$, and m is known, then $\mathcal{C}$ is non-parametrically identified from the transition measures $T(L_{jt}, L_{jt-1}, \pi_j)$.

Proof. Fix a starting debt principal L for a firm with expected profit π. From the transition measure, we can infer the probability of default for the given firm, $\mathcal{E}(L, \pi)$.

The interest rate set by the lender solves

$$\beta\big[(1+r)(1-\mathcal{E}(L,\pi)) + \delta\,\mathcal{E}(L,\pi)\big]L - L \;=\; mL.$$

Knowing r determines the next period debt liability, $L' = (1+r)\tilde{L}$. Once L' and π are known, uncertainty in the next period's liability comes entirely from next period shocks. This links the transition measure to the shock distribution:

$$T(L', L, \pi) \;=\; f_{\mathcal{C}}\big(L' + \pi - (1+r)L\big).$$

L' is bounded above by $L^{\max}(\pi)$. By continuity of r, we can find $\tilde{L}$ such that $(1+r)\tilde{L} - \pi = 0$ for every π. Call this $\tilde{L}^*(\pi)$.

Varying π, and starting from $\tilde{L}^*(\pi)$, we identify a region of support of $f_{\mathcal{C}}$:

$$T(L', \tilde{L}^*(\pi), \pi) = f_{\mathcal{C}}(L'),$$

which ranges over $[0, \tilde{L}^{\max}(\pi)]$. As $\pi \to \infty$, by assumption $L^{\max}(\pi) \to \infty$, which identifies the entire positive part of the distribution. If $\mathcal{C}$ is symmetric, it is fully identified.

Similarly, transitions $T(L^{\max}, 0, \pi)$ identify the negative part of the shock distribution, up to $\sup_\pi(L^{\max}(\pi) - \pi)$. Negative shocks beyond this range are not observed, since they move debt from the maximum level back to zero. $\square$

Intuitively, the idea is that by observing transitions starting from "just enough debt to break even," and then letting profits π grow arbitrarily large, we stretch out the observable support of shocks until the whole distribution is revealed. Symmetry then lets us recover the negative side.

Theorem 2.1. If m is known and we observe exit probabilities $\mathcal{E}(\tilde{L}, \pi)$ for profit types π from debt endowments $\tilde{L}$, and if $L^{\max}(\pi) \to \infty$ as $\pi \to \infty$, then $\mathcal{C}$ is non-parametrically identified.

Proof. Given $\mathcal{E}(\tilde{L}, \pi)$, the interest rate satisfies

$$\beta\big[(1+r)(1-\mathcal{E}(\tilde{L},\pi)) + \delta\,\mathcal{E}(\tilde{L},\pi)\big]\tilde{L} - \tilde{L} = m\tilde{L}.$$

Exit probabilities are related to the CDF of $\mathcal{C}$ by

$$1 - F_{\mathcal{C}}\big(L^{\max}(\pi) + \pi - \tilde{L}(1+r)\big) = \mathcal{E}(\tilde{L}, \pi).$$

Varying $\tilde{L}$, the derivative gives

$$f_{\mathcal{C}}\big(L^{\max}(\pi) + \pi - \tilde{L}(1+r)\big)\Big(r + \tilde{L}\frac{dr}{d\tilde{L}}\Big) \;=\; \frac{d}{d\tilde{L}}\,\mathcal{E}(\tilde{L}, \pi).$$

At $\tilde{L}^*(\pi)$ where $\tilde{L}^*(\pi)(1+r) = \pi$, we identify $f_{\mathcal{C}}(L^{\max})$, since $r(\tilde{L}, \pi)$ is known. Letting $\pi \to \infty$ (so $L^{\max}(\pi) \to \infty$) identifies the positive support of $\mathcal{C}$. Symmetry then gives the full distribution. $\square$

The intuition is that exit probabilities directly reveal the CDF of shocks at a shifted argument. By perturbing debt levels $\tilde{L}$, we shift this cutoff and trace out the density $f_{\mathcal{C}}$. Large profits π expand the observable range until the whole distribution is pinned down.

Theorem 2.2. If m is known and we observe exit probabilities $\mathcal{E}(L, \pi)$ for a fixed L across all profit types π, together with $L^{\max}(\pi)$, and if $L^{\max}(\pi) \to \infty$ as $\pi \to \infty$, then $\mathcal{C}$ is non-parametrically identified.

Proof. Given $\mathcal{E}(L,\pi)$, the interest rate solves

$$\beta\big[(1+r)(1-\mathcal{E}(L,\pi))+\delta\,\mathcal{E}(L,\pi)\big]L-L=mL.$$

Exit probabilities satisfy

$$1-F_{\mathcal{C}}\big(L^{\max}(\pi)+\pi-L(1+r)\big)=\mathcal{E}(L,\pi).$$

Differentiating in π,

$$-f_{\mathcal{C}}\big(L^{\max}(\pi)+\pi-L(1+r)\big)\left(\tfrac{d}{d\pi}L^{\max}(\pi)+1-L(1+\tfrac{dr}{d\pi})\right)=\tfrac{d}{d\pi}\,\mathcal{E}(L,\pi).$$

Starting from $L=0$ and letting $\pi\to\infty$, the assumption $L^{\max}(\pi)\to\infty$ identifies the whole distribution $\mathcal{C}$. □

The intuition here is parallel: varying π while holding L fixed changes the default margin, and observing how exit probabilities respond allows us to back out the shock density. Again, very high profits stretch the support to infinity, so the entire distribution is revealed.

Corollary 2. Any analytic parametric distribution of $\mathcal{C}$ is identified by observing either (a) $\mathcal{E}(L,\pi)$ for a set of (L,π) of positive measure, or (b) $(\mathcal{E}(L,\pi),L^{\max}(\pi))$ for a fixed L and a set of π of positive measure.

Proof. The arguments above show $f_{\mathcal{C}}$ is identified on a set S of positive measure. Choose $x\in S$ and a ball $B(x,r)\subset S$. By smoothness of $\mathcal{C}$, one can take derivatives of $f_{\mathcal{C}}$ at x. For an n-parameter family, these provide n moments to identify the distribution. □

Here the intuition is that once you can pin down the density on a set of positive measure, smoothness ensures that the whole parametric distribution is pinned down by matching derivatives (moments).

Optimal Planner Award Scheme

We propose a social planner's problem that allocates awards under different welfare weights on consumers and producers. Let the planner's per-period payoff be

$$R(s,a_i,a_{-i})\;=\;\lambda\,\mathrm{CS}(s,a_i,a_{-i})\;+\;(1-\lambda)\,\mathrm{PS}(s,a_i,a_{-i}),\qquad\lambda\in[0,1],\tag{38}$$

where CS and PS are consumer and producer surplus. Let $P(s)$ be the proposed policy instrument, which denote an award scheme that assigns 7(a) awards according to a probability distribution over states s. The social planner weighs welfare gains against the expected fiscal cost from guarantees $C(s, a_i, a_{-i}, P(s))$ in state $s = (\delta_i, \delta_{-i})$ given award status a_i, a_{-i} under scheme $P(s)$.

Holding the scheme P fixed, the continuation value from (s, a_i) satisfies

$$V(s, a_i, a_{-i}) \;=\; R(s, a_i, a_{-i}) \;+\; \beta \int_{s'} V(s', a_i', a_{-i}') Q(s' \mid s, a_i) \tag{39}$$

where the transition kernel $Q(\cdot \mid s, a_i)$ embeds how the award a_i shifts credit terms, default hazards, entry, and hence future market structure $(\delta_i', \delta_{-i}')$.

Equation (39) makes transparent that the value of an award has a *current* component via (38) and a *continuation* component via Q, i.e., by lowering exit risk of an awarded firm and altering future competitive structure. Lender maximizes $V(s, a_i, a_{-i}) - C(s, a_i)$. We describe the optimal solution to this planner's problem in detail below.

While the full dynamic optimization problem is computationally intractable due to both the high-dimensional state space and the complexity of value calculation, we can extract key economic insights from this formulation. [53] Efficient award allocation should focus resources where they generate the largest welfare gains per dollar of expected government recovery payments. Additionally, the value of supporting a firm depends not just on immediate effects, but also on whether that firm survives long enough to generate ongoing benefits through continued operation.

See now, the details of the optimal solution below.

Award scheme. Let a *cell* be indexed by incumbent-rival types (ϕ_i, ϕ_{-i}). A *scheme P* assigns a lifetime award at the time of entry:

$$a_i \sim \text{Bernoulli}(P(\phi_i, \phi_{-i}))\,, \qquad a_i \in \{0, 1\}.$$

Awards affect credit terms (limits, interest), default hazards, and hence entry, exit, and future market structure.

[53] Three factors make exact optimization infeasible. First, the state space of possible award configurations grows exponentially—with just 2 profit levels and 10 possible award levels, the number of award schemes exceeds 10^{16}. Second, even for a single award scheme, calculating values requires simulating many paths since the transition probabilities that determine payoffs are sparse and dispersed across states. Third, determining whether a scheme is feasible requires solving for equilibrium costs, which are themselves endogenous objects.

State, dynamics, and objective. Let s denote the industry state (incumbent types, award statuses, debt states, etc.). Given a scheme P, value satisfies the Bellman equation

$$V^P(s) \;=\; R_P(s) \;+\; \beta \sum_{s'} Q_P(s' \mid s)\, V^P(s'), \tag{40}$$

where Q_P is the law of motion induced by firms' entry/exit/default under P, and $R_P(s)$ is the planner's per-period payoff. Define the *discounted visitation (occupancy) measure*

$$\Lambda^P(s) \;:=\; (1-\beta) \sum_{t \geq 0} \beta^t \Pr(s_t = s \mid P),$$

so that the planner's objective can be written as

$$W(P) \;=\; \sum_s \Lambda^P(s)\, R_P(s).$$

Sensitivity (Gâteaux derivative). Consider a perturbation $P_\varepsilon = P + \varepsilon H$, where the direction H is concentrated on a single cell (ϕ_i, ϕ_{-i}). Differentiating (40) yields the standard sensitivity identity:

$$\frac{\mathrm{d}}{\mathrm{d}\varepsilon} W(P_\varepsilon)\Big|_{\varepsilon=0} \;=\; \sum_s \Lambda^P(s) \left\{ D_P R_P(s)[H] \;+\; \beta \sum_{s'} D_P Q_P(s' \mid s)[H]\, V^P(s') \right\}. \tag{41}$$

Decomposition into entry, award, and survival channels. Fix a current state s (with rival profile ϕ_{-i}), and a candidate entrant of type ϕ_i. Let $E_i(s; P)$ be the entry probability under P and let

$$\Delta W^{\text{cur}}(s; \phi_i) := \mathbb{E}\Big[R_P\big(s + \text{enter}_i^{a=1}\big) - R_P\big(s + \text{enter}_i^{a=0}\big) \Big]$$

be the *current* (within-period) payoff difference from awarding the entrant, conditional on entry. Define the entrant's next-period survival gain from the award, $\Delta S_i(s'|s) := S_i(s'|s, a_i{=}1) - S_i(s'|s, a_i{=}0)$, and the rivals' aggregate survival change, $\Delta S_{-i}(s'|s) := \sum_{j \neq i} \big[S_j(s'|s, a_i{=}1) - S_j(s'|s, a_i{=}0)\big]$. Let $\Omega_j^P(s') := \mathbb{E}\big[V^P(s' \mid j \text{ survives}) - V^P(s' \mid j \text{ exits})\big]$ denote the option value of survival for firm j. Then the continuation term in (41) can be written as

$$
\begin{aligned}
\beta \sum_{s'} D_P Q_P(s' \mid s)[H]\, V^P(s') = {} & \underbrace{\beta \, \frac{\partial E_i(s;P)}{\partial P(\phi_i, \phi_{-i})} \, \mathbb{E}\Big[V^P\big(s' \mid \text{enter}_i\big) - V^P\big(s' \mid \text{no enter}\big) \Big]}_{\text{(A) entry mass}} \\
& + \underbrace{\beta \, E_i(s;P) \, \mathbb{E}\Big[V^P\big(s' \mid a_i{=}1, \text{enter}_i\big) - V^P\big(s' \mid a_i{=}0, \text{enter}_i\big) \Big]}_{\text{(B) award mix at entry}} \\
& + \underbrace{\beta \, E_i(s;P) \, \mathbb{E}\big[\Delta S_i(s'|s) \cdot \Omega_i^P(s') \big]}_{\text{(C) entrant survival}} \\
& + \underbrace{\beta \, E_i(s;P) \, \mathbb{E}\big[\Delta S_{-i}(s'|s) \cdot \overline{\Omega}_{-i}^{P}(s') \big]}_{\text{(D) rivals' survival}},
\end{aligned}
$$

where $\overline{\Omega}_{-i}^{P}$ aggregates rivals' option values. Term (C) makes explicit that a lower future hazard (higher survival) for the award recipient raises the marginal value of awarding via Ω_i^P.

First-order condition (cellwise). With a cell-specific marginal expected cost $\kappa(\phi_i, \phi_{-i})$, the optimal scheme equates marginal benefit and marginal cost:

$$
\boxed{\sum_s \Lambda^P(s) \big\{ E_i(s;P) \cdot \Delta W^{\text{cur}}(s;\phi_i) \; + \; \text{(A)} + \text{(B)} + \text{(C)} + \text{(D)} \big\} \;=\; \kappa(\phi_i, \phi_{-i}) \, .}
$$

Under a global budget, replace κ with the common shadow price and obtain a bang–bang rule across cells.