NORTHWESTERN UNIVERSITY

Empirical Studies of Information, Market Power, and Policy Design

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

David Stillerman

EVANSTON, ILLINOIS

June 2022

Abstract

This thesis consists of three chapters: two are empirical studies of policy design in small-business lending markets, while the final is a large-scale retrospective of retail mergers.

Chapter 1 examines moral hazard in loan guarantee programs. To address credit constraints in small-business lending markets, policymakers frequently rely on loan guarantees, which provide lenders with insurance against default and aim to expand credit. Guarantees affect loan prices through two channels: (1) they alter the effective marginal cost of lending, and (2) they create a moral hazard problem, dampening incentives for lenders to collect information about borrower quality. The combination of these two effects implies that guarantee programs disproportionately benefit high-risk borrowers and may even harm low-risk borrowers. In this paper, I quantify these channels in the setting of the SBA 7(a) Loan Program and find that more generous loan guarantees lead to a decline in average loan prices, but that benefits are concentrated among high-risk borrowers. In fact, low-risk borrowers receive lower surplus as guarantees increase, and moving from the 90% guarantee observed in the data to a rate of 50% would raise their surplus by 2.5%. The heterogeneity in impact suggests that alternative policies that moderate the effect of bank moral hazard could increase aggregate borrower surplus. I propose a hybrid policy design with a 50% guarantee and a subsidy set such that government spending is fixed. When compared to the baseline of 90% with no subsidy, the alternative policy mitigates the redistribution from low- to high-risk borrowers and leads to a 1.6% increase in borrower surplus, on average, and 0.1 percentage point (1.6%) decline in the program's default rate.

Chapter 2, which is joint work with Paul Kim, studies the incentive design of the Paycheck Protection Program, a large loan forgiveness program implemented during the Coronavirus pandemic. The program was executed through private lenders with the goal of assisting small businesses in keeping their employees on payroll during the COVID-19 pandemic.

We develop a model of PPP lending to capture the government's tradeoff between inducing bank participation and targeting funds for use on payroll. Using the model, we establish that both increasing subsidies and relaxing forgiveness standards are effective in expanding credit access to borrowers seeking smaller loans. However, their efficacy in targeting (i.e., providing funds to businesses who will use them on payroll) depends on the correlation between loan amounts and borrowers' return to payroll. We test the implications of the model using policy variation from the PPP Flexibility Act, legislation that relaxed forgiveness standards. Consistent with the predictions of the model, the average loan amount falls by between 6 and 7% in the period following the policy change. Furthermore, marginal borrowers are more likely than inframarginal borrowers to use funds for payroll, so making forgiveness more accessible increases the average share of funds used for those purposes.

Chapter 3, which is joint work with Vivek Bhattacharya and Gastón Illanes, is a large-scale retrospective of US retail mergers. In particular, we document the price and quantity effects of all US retail mergers from 2006–2017 associated with deals larger than $340 million. Prices increase by 0.49% on average for merging parties, with an interquartile range of almost 5%. Non-merging parties exhibit slightly smaller price changes on average. Total quantities decline on average by 3–5%, but there is even larger variation across mergers. We investigate the role of synergies and market power by analyzing the timing of price changes, the relationship with locations of production facilities, and measures of market structure. We collect data on merger enforcement (remedy proposals in our case), and through the lens of a simple model, we estimate that agency preferences are such that they aim to challenge mergers where prices are expected to increase by more than 3.7–5.6% overall, or about 8.1–8.8% for merging parties.

Acknowledgements

Thank you to my committee for their support throughout the development of this work. First, I want to thank my chair, Robert Porter, for listening as I pitched ideas, worked through problems with estimation, and wrote this thesis. Especially as the pandemic took us away from campus for two years, I looked forward to our weekly meetings, and they helped guide this work to where it is today. Second, thank you to Vivek Bhattacharya and Gastón Illanes for taking a chance on me as an RA during my second year. This piqued my interest in the intersection of IO and finance, and I can confidently say this thesis would not be where it is without Vivek and Gastón. They have been incredibly generous – sometimes, too generous – with their time. I feel fortunate to have had them as mentors and to be working with them as coauthors on the third chapter of this thesis.

I would also like to thank a number of Northwestern faculty members for helpful discussions, comments, and suggestions, including Gregor Matvos, Mar Reguant, Bill Rogerson, and Paola Sapienza. A number of students in the department also had an instrumental impact on my graduate school experience and on the development of this work. In particular, thank you to Sam Goldberg, Joao Guerreiro, Paul Kim, Kristina Manysheva, Joris Mueller, and Richard Peck.

Finally, thank you to my friends and family for for being here every step of the way and keeping me grounded during these past six years. In particular, thank you to my mom for welcoming me home to Evanston with open arms and to Madison and Adelaide for their unwavering support (and patience) as this thesis has come together. I could not have done it without you all.

Table of Contents

List of Tables

List of Figures

Chapter 1

Loan Guarantees and Incentives for Information Acquisition

1.1 Introduction

Small businesses are a major source of employment growth in the United States. They are responsible for a sizable share of private-sector jobs, employing 47.3 percent of the workforce across 30.7 million small businesses as of 2016.[1] Given their large contribution to aggregate employment and the substantial attention paid to them by the media and general public, governments tailor policy to foster small-business growth. One common piece of the policy agenda, both in the United States and elsewhere, involves expanding access to capital. Many small businesses face large up-front costs, including real estate and machinery, and traditionally find it difficult to obtain start-up funding and capital required for ongoing operational expenses. Since small firms are typically unable to access institutional debt markets or equity financing (Mills and McCarthy (2014)) and rely disproportionately on traditional bank

[1]"2019 Small Business Profile." U.S. Small Business Administration Office of Advocacy. Available at https://cdn.advocacy.sba.gov/wp-content/uploads/2019/04/23142719/2019-Small-Business-Profiles-US.pdf.

lending, governments typically transmit small-business policy through this channel.

Access to credit is further constrained by information asymmetries, which are prevalent in lending markets. Recent empirical work notes their importance in consumer lending (see, e.g., Einav et al. (2012); Cuesta and Sepulveda (2021)), mortgages (Allen et al. (2019)), and commercial lending (Crawford et al. (2018); Ioannidou et al. (2022); Wang (2020)). In small-business credit markets, the degree of asymmetric information is potentially larger. Nascent businesses lack a long repayment history and a standardized credit score, like the FICO score used in consumer lending, from which lenders can judge a borrower's creditworthiness. Thus, lenders leave many decisions up to the discretion of loan officers and delegate much of the information acquisition and loan issuance process to these individuals. Much of the time, officers rely on "soft" information[2] (e.g., meeting a borrower, judging the borrower's trustworthiness, verifying revenue projections, etc.) when deciding whether to approve a loan and what interest rate to offer (Wang (2020)).

Seminal work by Stiglitz and Weiss (1981) shows that asymmetric information in lending markets can lead to credit rationing and price-risk mismatch. Both results contribute to credit constraints faced by small businesses and induce policymakers to address these market failures. One common intervention, and one that is typically prioritized during an economic downturn, is a loan guarantee program. Guarantees act as insurance against ex-post default risk: governments agree to cover a pre-specified portion of the remaining loan balance in the event of default. Such an intervention lowers the cost of default faced by the lender and provides incentives for them to extend funds to riskier borrowers. However, guarantees also create a moral hazard problem by weakening lenders' incentives to properly evaluate borrower risk.

[2] See Liberti and Petersen (2019) for a discussion of "hard" vs. "soft" information. The key distinguishing factor is the importance of context or subjective opinion. "Hard" information typically can be described numerically and its interpretation must not depend on subjective judgment. "Soft" information instead can be interpreted differently based upon who collects it or how it is collected.

The aggregate impact of a guarantee program is driven by the two channels described above. First, guarantees alter the effective marginal cost of lending. Default is less costly to the lender when guarantees are generous, exerting downward pressure on prices. In addition, in the presence of adverse selection, guarantees decrease the difference between the cost of serving the marginal borrower and that of serving inframarginal borrowers. In a market characterized by adverse selection, the marginal borrower is less risky than inframarginal borrowers, and this constrains a lender's willingness to raise prices (Mahoney and Weyl (2017); Crawford et al. (2018)). Guarantees alter this disincentive, exerting upward pressure on prices, and could offset the gains that would otherwise accrue to borrowers due to the lower cost of default. For the duration of the paper, I refer to the price effects resulting from the change in effective marginal cost as the *guarantee pass-through*, and its magnitude depends on the probability of default, the extent of adverse selection in the market (i.e., the correlation between price responsiveness and propensity to repay), and borrowers' willingness to pay for funds.

Second, the guarantees affect lenders' incentives to acquire information and, in turn, create a moral hazard problem. They bring the lenders' payoffs in the two states of the world, repayment and default, closer to one another. Thus, increasing guarantees decreases the marginal value of additional information. In simple terms, the lender now gains less by being able to distinguish between a high- and low-risk borrower and, assuming information is costly to acquire, could lead to an equilibrium with information that is less precise.[3] I refer to this as the *information effect*. This unintended consequence is welfare-improving for high-risk borrowers but harms those that are low risk, which benefit from lenders having precise information. Policymakers have debated the potential for moral hazard in loan guarantee programs, most recently in the United States Senate Budget Committee in 2014[4], but there

[3]This idea is similar to the "lazy bank hypothesis" set forth by Manove et al. (2001) in their examination of the effect of collateral on lender behavior. However, adjusting the government guarantee rate does not affect borrower incentives, whereas collateral mitigates both adverse selection and ex-post moral hazard.

[4]The Ranking Member of the Committee writes: "The guarantee takes a substantive risk away from the

is no consensus on the magnitude of these effects. This paper provides the first estimate of this magnitude.

The combination of the two channels described above implies that an increase in the generosity of guarantees need not benefit all borrowers. Determining the direction and magnitude of the impact of guarantee programs is an empirical question, and alternative policy arrangements may lead to higher borrower surplus. In this paper, I answer three questions. First, does an increase in the generosity of guarantees lead to higher borrower surplus, on average? Second, is the increase in borrower surplus a Pareto improvement? Third, does an alternative policy design lead to better outcomes when judged on a borrower-surplus standard? Understanding the separate impact of the guarantee pass-through and the information effect is crucial for evaluating the aggregate effect of alternative policy designs, as well as their distributional impact. In particular, the strength of the information effect influences how responsive lenders are to changes in their risk exposure. For example, a strong information effect suggests that a policy that leaves lenders more exposed to default risk may lead to gains in borrower surplus.

I answer these questions using data from the U.S. Small Business Administration's (SBA) 7(a) Loan Program, which is the largest lending program administered by the agency and provides guarantees on loans to small businesses. Between 2009 and 2010, Congress enacted two separate pieces of legislation which contained components aimed at stimulating small-business lending. The American Recovery and Reinvestment Act of 2009 and the Small Business Jobs Act of 2010 allocated a total of \$1.2 billion (\$730 million in the former and \$510 million in the latter) to increase maximum guarantee rates and reduce guarantee fees.[5]

bank making the loan (indeed, that is the whole point of the SBA loan guarantee) and very well could provide incentives for some banks to cut corners through the underwriting process. This is what economists call moral hazard and could have manifested itself with lenders being less than careful in their decisions to extend SBA loans." For more details, see https://www.budget.senate.gov/chairman/newsroom/press/sessions-writes-to-small-business-administration-about-cost-of-loan-program.

[5]Congressional Research Service (2019b), available at https://crsreports.congress.gov/product/pdf/RL/RL33243.

This resulted in two discrete time periods during which government guarantees were higher. Using the temporal variation induced by these policy changes, I provide evidence of a strong lender-side response to the guarantee incentives.

Higher guarantees induce both a spike and sustained increase in lending activity, and lenders offer contracts with more generous terms: lower interest rates, larger amounts, and longer maturities. In the guarantee expansion (SBA Recovery) period, average interest rates are 4.2 basis points lower than in the baseline (1% of the mean rate), average loan amounts are $56 thousand higher (10% of the mean amount), and average loan terms are between 4 and 5 months longer (3% of the mean term), all economically and statistically significant changes. The guarantee-rate increase also leads to a lower correlation between loan prices and ex-post default, indicating a decline in the precision of risk pricing when guarantees become more generous. Furthermore, the lower correlation is driven by large price declines for high-risk borrowers. In fact, borrowers at the lower tail of the risk distribution, as measured by expected default, receive higher price offers in the SBA Recovery period than in the baseline. These results imply that, in equilibrium, guarantees have a heterogeneous impact across the risk distribution.

The descriptive evidence by itself illustrates the aggregate effect of loan guarantees, but it is unfit to examine the mispricing of risk and its effect on borrower surplus. Moreover, it does not allow for an examination of alternative policy designs. For these purposes, I develop a model of small-business lending with guarantees. Borrowers are differentiated by observable characteristics, as well as ex-ante unobservable risk and price responsiveness. Lenders choose the precision of information to acquire, given the known distribution of borrowers, and receive noisy signals of risk. The precision of these signals governs how closely price offers match underlying borrower risk. Under the structure of the model, I recover the joint distribution of the borrowers' risk and price responsiveness, as well as the primitives governing their default and acceptance decisions. I then describe how the difference in the distribution of price offers conditional on default and that conditional on repayment pins down the information effect. The relative magnitude of this information effect and the guarantee pass-through determine the extent to which the equilibrium outcomes of the policy disproportionately affect high-risk borrowers.

I estimate the primitives of the model using loan-level data provided by the SBA. The estimates imply that banks respond to the decline in risk by collecting noisier signals of borrower quality. On average, the standard deviation of the noise in the lenders' signals is 7% higher in the guarantee expansion period, though the effect is heterogeneous across lender types. On the borrower side, there is substantial observable heterogeneity in borrowers' propensities to default and utilities of loan acceptance. These quantities vary across types of borrowers and across time. For example, loans issued in 2010 and 2011, during the recovery, are associated with lower likelihoods of default than loans issued at the beginning of the recession in 2009. Unobservable risk and price responsiveness are positively correlated, indicating the presence of adverse selection.

With these estimates in hand, I examine equilibrium price offers, and the corresponding borrower surplus, across counterfactual guarantee rates to disentangle the impact of the guarantee pass-through and of adjustments to lenders' screening (information effect). This analysis quantifies the effect of the lenders' weaker incentives to gather information on borrower outcomes and illustrates how it magnifies the distributional impact of an increase in the guarantee rate. For guarantee rates between 50% and 100%, prices decrease and borrower surplus increases for the average borrower as guarantees become more generous. The guarantee pass-through and information effect work in the same direction. Holding the precision of information fixed, prices are 0.1% higher and borrower surplus 0.4% lower under a guarantee of 50% than under a guarantee of 90%, which is the rate observed in the SBA Recovery period. Allowing for flexible information acquisition magnifies these effects, adding an additional 0.1% increase in price and 0.6% decline in borrower surplus. However, borrowers face starkly different outcomes depending on their risk. In total, low-risk borrowers (i.e., those in the top quintile of the distribution of the utility of repayment) receive 2.5% more surplus under the guarantee of 50% than in the baseline, while their high-risk counterparts (i.e., the bottom quintile of the distribution of the utility of repayment) receive 6.6% less.

This decomposition illustrates that both the guarantee pass-through and information effect disproportionately benefit high-risk borrowers relative to their peers. If the intent of the program is to expand credit to the most credit-constrained borrowers, then the equilibrium outcomes suggest the program is successful in that aim. That being said, the gains to high-risk borrowers come at the expense of lower-risk borrowers in the 7(a) program. This latter point suggests that alternative policies that limit the impact of bank moral hazard could lead to better outcomes for low-risk borrowers and potentially lead to aggregate surplus gains.

A natural choice is a hybrid policy with a less generous guarantee and a subsidy. The less generous guarantee provides stronger incentives for lenders to acquire information and offsets some of the disproportionate gains accruing to high-risk borrowers. The addition of the subsidy decreases the cost of lending for all borrowers, inducing downward pricing pressure. I find that a program with a guarantee of 50% and a subsidy such that expected spending is fixed generates a 1.6% increase in borrower surplus, on average, over the baseline guarantee of 90% with no subsidy. This policy also reduces the heterogeneity in the impact across the distribution of borrower risk, raising surplus for low-risk borrowers and decreasing it for their high-risk peers. The resulting change in the composition of borrowers that accept loans leads to a 0.1 percentage point (1.6%) decline in the program's default rate. These findings suggest that a combination of policy instruments can be used to reduce the impact of bank moral hazard and increase borrower surplus, while limiting default risk.

The remainder of the paper proceeds as follows. Section 1.2 reviews the relevant literature, Section 1.3 presents a stylized model of small-business lending and illustrates comparative statics results, Section 1.4 describes the SBA 7(a) Loan Program and policy changes, Section 1.5 describes the SBA dataset and presents descriptive analyses, Section 1.6 introduces the model and discusses estimation of the primitives, Section 1.7 discusses the estimates, Section 1.8 presents results of the decomposition exercise and counterfactual policy simulations, and Section 1.9 concludes.

1.2 Related Literature

This paper builds on three main strands of literature. First, it contributes to the empirical analysis of asymmetric information and contract design. The work in this literature has examined many different markets, yet it is all built around one key idea. It leverages the fact that ex-post outcomes are observable in many contexts (e.g., claims in health insurance, accidents in auto insurance, or default in lending), which allows the econometrician to analyze selection into contracts with particular characteristics (e.g., price, down payment, etc.) by consumers who are observably different ex post. Seminal work by Chiappori and Salanie (2000) develops a test for asymmetric information using auto insurance plan selection, and similar ideas have recently been applied to other markets with information asymmetries, including health insurance[6] and lending markets. Einav et al. (2012) examine these ideas in the context of subprime auto loans, quantifying the effect of contract characteristics on the types of borrowers that choose a given contract and the resulting default rates. They create a novel framework that allows for flexible correlation between consumer willingness to pay and ex-post loan performance (i.e., repayment). Much of the recent work examining asymmetric information in lending markets follows a similar structure. For example, Crawford et al. (2018) model the interaction between imperfect competition and asymmetric information in the commercial lending market and find that markets in which both are present face lower

[6]For example, Starc (2014) models contract design for Medigap insurance plans. In the model, insurers account for endogenous changes in healthcare utilization stemming from changes to contract characteristics. In the lending context, similar ideas are present and motivate my decision to allow lenders to account for selection on unobserved consumer type.

welfare losses than markets in which only one is present. Yannelis and Zhang (2021) examine this same idea, but allow for endogenous information acquisition when screening costs are fixed. My paper builds upon this strand of work that analyzes the endogenous collection of information, in this case focusing on potential moral-hazard effects of government guarantee policies in a market with variable screening costs. I incorporate ideas from the theory literature (e.g., Stiglitz and Weiss (1981) and Manove et al. (2001)) and empirical literature on screening soft information (e.g., Panetta et al. (2009) and Wang (2020)) to microfound the information structure of the contract designer and apply this framework to analyze loan guarantee systems.

Second, this paper contributes to the literature on bank moral hazard. A large theoretical literature examines bank moral hazard and incentives for ex-ante screening (see, e.g., Gorton and Pennacchi (1995)) and ex-post monitoring (see, e.g., Holmstrom and Tirole (1997)), but empirical work is limited. The majority of empirical papers analyze moral hazard in the context of loan securitization and exploit a rule of thumb that induces a higher likelihood of securitization above a credit-score threshold. Keys et al. (2010) show that securitization leads lenders to exert less scrutiny over potential borrowers, as securitized loans default at a higher rate than observably similar, unsecured loans. Rajan et al. (2015) exploit the same variation but instead examine loan characteristics. To the best of my knowledge, Rajan et al. (2015) provide the first comprehensive analysis of the impact of bank moral hazard on the mapping from loan characteristics (e.g., price) to outcomes and find that greater likelihood of securitization leads to a noisier mapping from price to default. One additional paper examines similar ideas outside the securitization setting. Avramidis et al. (2021) examine screening incentives in small-business lending, exploiting policy variation in the Community Reinvestment Act to show that lenders collect less information in competitive markets. My paper builds on the ideas of these three empirical papers but focuses on the heterogeneous impact of bank moral hazard across the distribution of borrowers. The ideas underlying the distributional impact of moral hazard are related to the findings of Nelson (2022), who studies the impact of an increase in pooled pricing, precipitated by the 2009 CARD Act's limits on lenders' abilities to raise borrowers' credit card interest rates over time, on borrowers across the distribution of risk. Assessing the distributional impact of government policy is particularly notable in the context of loan guarantee programs, as the policies are designed with the goal of expanding credit to riskier borrowers that have limited outside funding options.

Finally, this paper relates closely to work examining the efficacy of guarantee schemes. Bachas et al. (2021) exploit the kink design of the SBA 7(a) Loan Program's guarantee schedule to estimate the elasticity of credit supply with respect to the guarantee rate. Leveraging similar methods as in the literature on bunching estimators, they find that a one-percent increase in the guarantee rate leads to \$19,000 of additional lending. Cox et al. (2021) examine market power in the SBA's Express Loan Program and find that the guarantee program has a small, positive impact on both borrower and lender surplus. I find similar, small average effects of such a policy, but in this paper I stress that examining aggregate effects may mask heterogeneity across the risk distribution. Ioannidou et al. (2018) study how policy variation in guarantee rates changes lenders' incentives to collect collateral. The incentives described by these authors are similar to those underlying information acquisition decisions explored in my paper, but the effects of bank moral hazard on pricing differ from outcomes associated with changes in collateral.

Other recent work examines the aggregate effects of loan guarantees. Brown and Earle (2017) estimate the impact of expanded credit access on small-business employment growth. They rely on data from the SBA 7(a) guarantee program and find that increased access to SBA-affiliated lenders leads to employment expansion. They take this as evidence that, absent the guarantee program, small businesses face noticeable credit constraints. Choi and Lee (2019) exploit the same policy shock as I do and find that smaller banks respond sharply to high guarantee rates by expanding lending activity but larger banks do not. While these papers examine aggregate responses of both banks and small businesses to guarantee schemes, my paper instead focuses on individual loan outcomes and pricing, explicitly accounting for selection on unobservable borrower risk.

1.3 Simplified Model

To set ideas and illustrate the incentives at play in lending markets with guarantee programs, I now present a stylized model of lending interactions. This example demonstrates the competing forces at play in this setting and highlights the heterogeneous effect, across the distribution of borrower risk, of loan guarantees on prices offered to borrowers. I highlight the primitives on both the borrower and lender side that determine the sign and magnitude of the aggregate effect of a guarantee-rate change on borrower surplus. The empirical model, which I introduce in Section 1.6, builds on these same primitives.

The results in this section are driven by the two effects described in Section 1.1: the *guarantee pass-through* and the *information effect*, an endogenous change in the lender's information quality due to moral hazard. Absent any adjustments to the precision of information, an increase in the guarantee rate induces a change in price for all borrowers that default with positive probability, but the magnitude of the change varies with the level of risk. In this example, an increase in guarantee generosity leads to a price decline, and the magnitude of the effect is decreasing in the borrower's perceived ability to repay. However, as discussed in Section 1.1, guarantees need not lead to lower prices for all borrowers in a more general framework, such as the model I present in Section 1.6.[7] In a market with adverse

[7] This result is related to the empirical analysis of Crawford et al. (2018) and the theoretical work of Mahoney and Weyl (2017), both of which examine the interaction between adverse selection and market power. Their main finding is that market power mutes the effects of adverse selection and vice-versa. In particular, Crawford et al. (2018) distinguish between the effect of price changes on the relative default rates of the average and the marginal borrower, the latter of which underlies the upward pricing pressure stemming from an increase in the generosity of guarantees.

selection, the positive relationship between willingness to pay and default constrains the lender's willingness to raise prices. The guarantee dampens this disincentive of a marginal price increase, which exerts upward pressure on prices. A change in the precision of the lender's information has a similar, heterogeneous impact across borrowers. In total, the aggregate effect on price offers, and on borrower surplus, depends on the relative strength, and direction, of the guarantee pass-through and the effect of a change in the lender's information precision.

Suppose a borrower seeks a loan of a pre-determined size from a lender that has market power. Further, suppose the borrower takes one of two types, high risk (ξ^H) or low risk (ξ^L), and high-risk borrowers have (1) a higher willingness to pay and (2) a higher probability of default than low-risk borrowers. The lender is fully informed of the distribution of borrower types in the population, but not of the individual borrower's type, and instead receives a signal of that borrower's risk. The lender then updates its beliefs of the borrower's type and presents a price offer. That price offer is a function of default risk, the loan guarantee rate, and the borrower's willingness to pay. The lender maximizes its expected payoff, conditional on the signal it receives of the borrower's type:

$$p^* = \text{argmax}_p \sum_{\xi \in \{\xi^H, \xi^L\}} Pr(\xi|s) \left[P^A(p, \xi) \left[(1 - (1 - M) P^D(\xi)) p - \zeta \right] \right],$$

where p is the loan price, $P^A(\cdot, \cdot)$ is the mapping from price and borrower type to the probability of acceptance, $P^D(\cdot)$ is the mapping from borrower type to default, s is a binary signal of borrower type, M is the guarantee rate, and ζ is the marginal cost of lending. The fact that the borrower may default and the lender may be reimbursed in accordance with the guarantee rate acts as an adjustment to the cost of lending. Reformulating the problem in terms of acceptance probabilities (i.e., "quantities" demanded by a single borrower), the effective marginal cost is given by:

$$MC(P^A) = \sum_{\xi \in \{\xi^H, \xi^L\}} Pr(\xi|s) \left[\zeta + (1-M) P^D(\xi) \left(p(P^A, \xi) + P^A \frac{\partial p}{\partial P^A}(P^A, \xi) \right) \right].$$

The guarantee pass-through is evident in the structural relationship between the guarantee rate, M, and the effective marginal cost. The magnitude of this effect depends on: (1) the default probability and (2) the price responsiveness of each borrower type. Because a higher guarantee rate decreases the lender's losses in the case of default, price offers respond more strongly to guarantee-rate changes when default outcomes are likely. Thus, a larger share of the guarantee is passed through to high-risk borrowers than to their low-risk counterparts. Borrowers' price responsiveness also determines the extent of the pass-through. If acceptance is more responsive to a change in price, lenders are less able to pass through any change to the guarantee rate. The information effect instead enters through the lender's posterior beliefs of the distribution of borrower types. The magnitude of this effect depends on how strongly lenders adjust their precision of information in response to policy changes, as well as how heterogeneous borrowers are in the population.

In the subsections below, I examine two sets of comparative statics. First, in Section 1.3.1, I isolate the direct effect of changes to the guarantee rate by assuming that lenders are unable to adjust their information quality in response to the new government policy. I show that the magnitude of this effect depends on borrowers' price responsiveness. Then, in Section 1.3.2, I examine the effect of a change in signal precision, holding the guarantee rate constant. In sum, these results illustrate the ambiguous effect of guarantee-rate changes on the equilibrium prices offered to borrowers and describe which primitives determine the magnitude of the aggregate effect on prices and, thus, borrower surplus.

1.3.1 Guarantee-Rate Change

Figure 1.1 illustrates the effect of an increase in the guarantee rate, all else equal. For simplicity, both before and after the policy change, assume that there is an equal proportion of high and low types in the population, and that the lender receives an uninformative signal of the borrower's type. Also, assume that demand is linear in the price offer and that the probability of default is fixed for each borrower type with $P^D(\xi^H) > P^D(\xi^L)$. Panel (a) displays the demand schedules for borrowers of each type, and risk and willingness to pay are positively correlated. Panel (b) unpacks the lender's pricing decision, and displays expected demand, expected marginal revenue, and expected marginal cost under the low- and high-guarantee policies.

An increase in the guarantee rate rotates the expected marginal cost leftward. This rotation is driven by two underlying cost components. First, a higher guarantee rate makes default less costly to the lender, decreasing the cost of serving all inframarginal borrowers. Second, a more generous guarantee decreases the difference between the cost of serving the marginal borrower and that of serving the average borrower. Together, these two channels imply that higher guarantees decrease the cost associated with default, flattening the mapping from acceptance probability to effective marginal cost. In the limit, the cost approaches a constant, ζ, for all acceptance probabilities.

As mentioned previously, in this example the leftward rotation in costs results in lower prices for all borrowers. In equilibrium, borrowers of both types accept loans with higher probability and obtain higher borrower surplus than they do under lower guarantee rates. However, the magnitude of the price and borrower-surplus changes depends on each individual borrower's willingness to pay, their probability of default, the extent of heterogeneity in the population, and the lender's baseline marginal cost. Recovering each of these primitives is a requisite for quantifying the pass-through of a guarantee policy change and is one goal of the remainder of this paper.

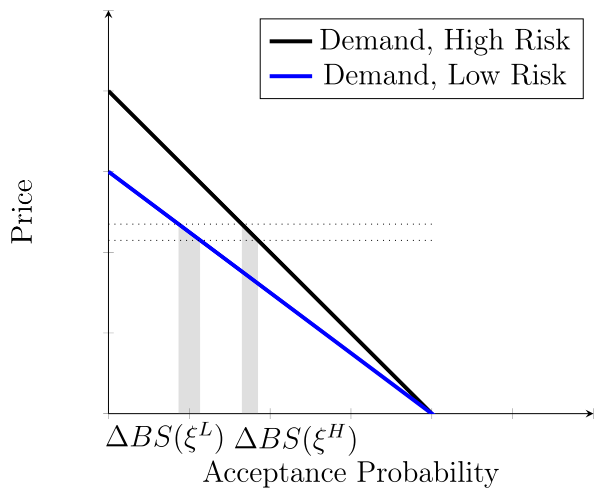

(a) Demand Schedules

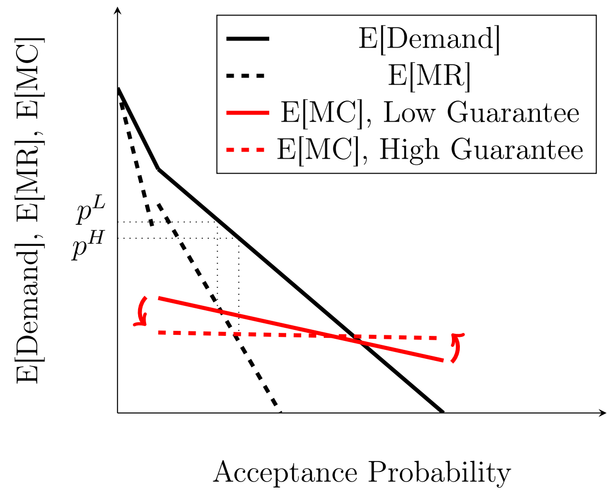

(b) Expected Demand, MR, and MC

Figure 1.1: Guarantee-Rate Change

Note: $\Delta BS(\tilde{\xi})$ denotes the change in surplus for a borrower of type $\tilde{\xi}$. Gray shading indicates an increase in borrower surplus when the guarantee becomes more generous, while red shading indicates a decline.

The declining relationship between the magnitude of price and borrower-surplus changes and the borrower's ability to repay relies on the assumption that the precision of information is held constant. As I demonstrate in Section 1.3.2, allowing for flexible screening yields non-monotonic price changes and results in an equilibrium shift toward riskier borrowers, at the expense of serving those that are less risky.

1.3.2 Information Structure Change

Figure 1.2 depicts the pricing response, and the subsequent change in acceptance rates and borrower surplus across types, to a decline in the lender's precision of information. The demand schedules for high- and low-risk borrowers are identical to those in the previous section, but the lender faces different expected demand depending on the realized signal and the precision of information. Suppose, for ease of exposition, that the lender receives a binary signal indicating the borrower is a high type. Consider the pricing response under two assumptions about signal precision. First, displayed in black on Panel (b) are expected demand, marginal revenue, and marginal cost for a lender that receives an informative signal (i.e., a signal, s, such that $Pr(\xi^H|s) > Pr(\xi^H)$) that the borrower is a high-risk type. The red figures illustrate those same three quantities for a lender that receives an uninformative signal (i.e., a signal, s, such that $Pr(\xi^H|s) = Pr(\xi^H)$). Panel (d) repeats the analysis in Panel (b) under the assumption that the lender receives a signal that the borrower is a low type, and the curves corresponding to the uninformative signal are again shown in red.

When the lender's information becomes less precise, high-risk borrowers receive lower price offers, while low-risk borrowers receive the opposite. High-risk borrowers benefit from noisier signals (i.e., a riskier borrower prefers that the lender has a more difficult time distinguishing it from a less risky borrower), while low-risk borrowers instead face a decline in borrower surplus. Combining this result with the discussion in Section 1.3.1, a change to the guarantee rate has an ambiguous effect on the price offer received by an individual borrower. Importantly, the direction of the price-offer change depends on the relative strength of the guarantee pass-through on the lender's effective marginal cost and the effect of an endogenous change in the lender's precision of information. For low-risk borrowers, the indirect effect could offset, or even reverse, the direct effect on prices. On the other hand, for high-risk borrowers, the information effect could amplify price declines.

For the remainder of this paper, I focus on quantifying the primitives governing the borrowers' acceptance and default decisions, as well as the distribution of risk. These primitives determine whether a more generous guarantee policy increases aggregate borrower surplus, as well as the magnitude of the change. I then examine the lender's pricing decision and develop a strategy to quantify changes in the lender's precision of information separately from direct effects of loan guarantees. Using these estimates, I decompose the reduced-form effects of a change to guarantees into the direct contribution of the ex-post insurance and the contribution of endogenous changes to the quality of the lenders' information. Furthermore, I then examine counterfactual guarantee policies and their impact on equilibrium price offers

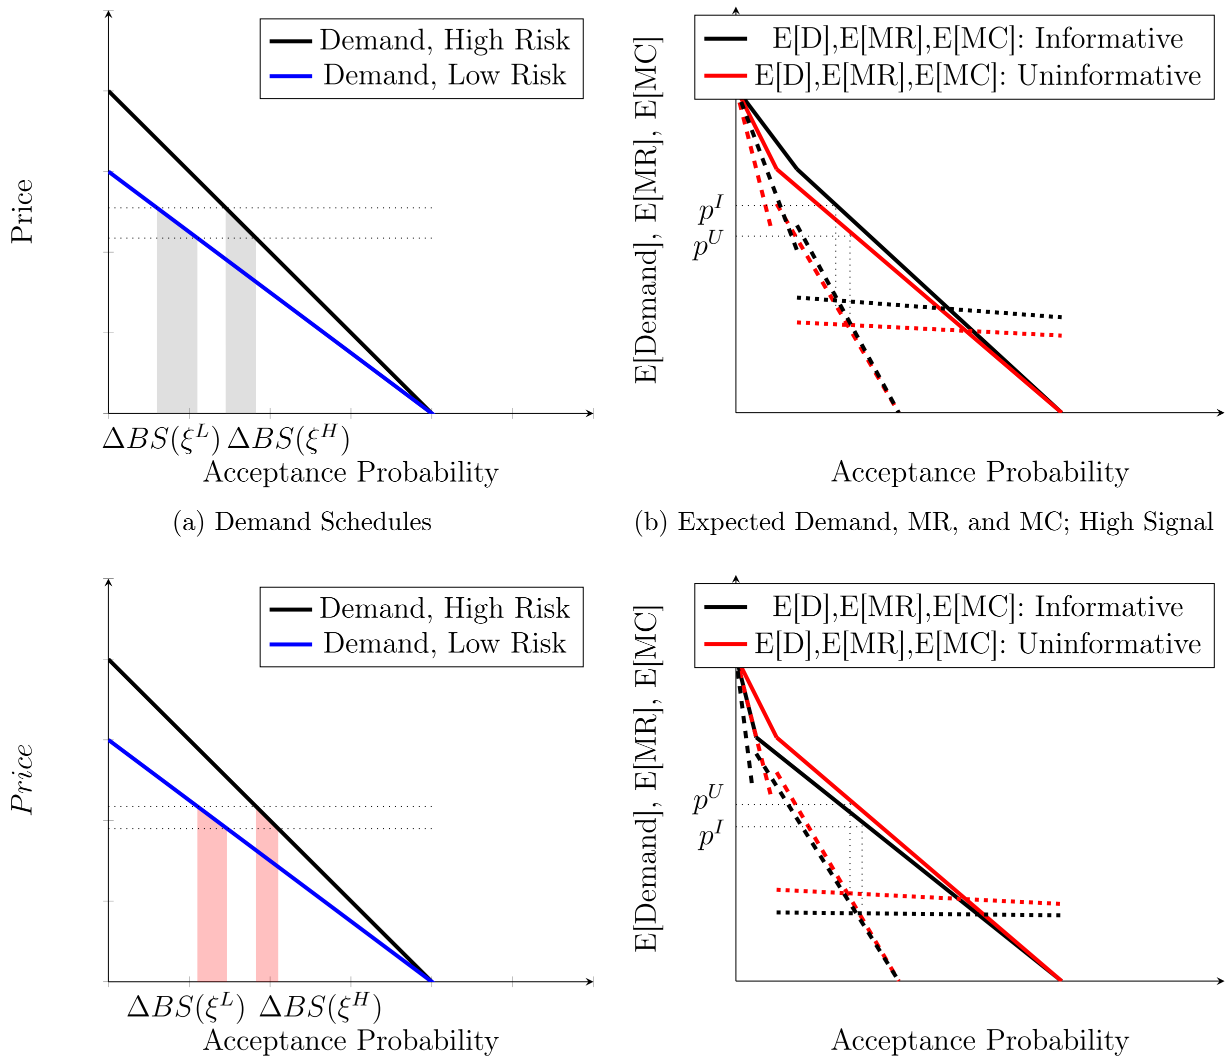

(a) Demand Schedules

(b) Expected Demand, MR, and MC; High Signal

(c) Demand Schedules

(d) Expected Demand, MR, and MC; Low Signal

Figure 1.2: Information Structure Change

Note: $\Delta BS(\tilde{\xi})$ denotes the change in surplus for a borrower of type $\tilde{\xi}$. Gray shading indicates an increase in borrower surplus as information becomes less precise, while red shading indicates a decline.

to borrowers across the risk distribution.

1.4 Institutional Background

The focus of the empirical analysis in this paper is the SBA 7(a) Loan Program (now known as the SBA Advantage Loan Program), which is the SBA's main loan guarantee program and was established by the Small Business Act of 1953. The program's principal intent is to provide financing to small businesses, particularly those that are credit-constrained and have few outside funding options. In the words of the SBA, these businesses have "sufficient cash flow to repay the loan but may not have the necessary collateral or history required by a bank's lending policy."[8] To be eligible for a loan issued under the 7(a) program, businesses must satisfy size standards set by the SBA and pass a *credit elsewhere test*. While the credit-elsewhere rule does not rely on a quantifiable metric, it states that SBA guarantees are available only to borrowers that would be unable to obtain a loan absent the guarantee.[9] This stipulation motivates my decision to model lenders with monopoly power in Section 1.6. Competition among lenders is limited to those within the SBA 7(a) program and, typically, a borrower applies only to a single lender. Only 282 of 51,153 borrowers that were approved for an SBA guarantee between 2009 and 2011 were approved with more than one lender during that same period.

Funds obtained through 7(a) loans must be used for an allowed purpose. For example, they may be used to purchase real estate or manufacturing equipment, to make capital improvements, or to acquire or start businesses, among other uses. However, borrowers cannot take out additional SBA-backed funds to refinance an existing loan in the 7(a) program.[10]

[8] Office of the Comptroller of the Currency (2015), available at https://www.occ.gov/publications-and-resources/publications/community-affairs/community-developments-fact-sheets/ca-fact-sheet-what-is-sba-7a-loan-guaranty-jul-2015.html.

[9] See 13 CFR 120.101 and 13 CFR 120.102. The credit elsewhere test requires that the lender review available resources of any individual who owns at least a 20% stake in the firm and must identify a deficiency in the borrower's profile, such as lack of sufficient collateral, that renders the borrower unable to receive funds absent a guarantee. More information on this requirement can be found at: https://www.sba.gov/offices/district/mt/helena/resources/lenders-8-first-steps-determine-sba-eligibility-and-prevent-application-processing-delays.

[10] Office of the Comptroller of the Currency (2015)

These funds reach a sizeable number of borrowers each year. For instance, in the fiscal year 2018, the SBA guaranteed 60,353 loans, worth approximately $25.4 billion.[11] While only a small subset of the 30 million small businesses in the U.S. receive loans through the program, these firms tend to be particularly credit constrained, as evidenced by their passage of the credit elsewhere test.

The lending process for guaranteed loans is structured. It consists of a number of steps, though these steps vary depending on whether the lender is a participant in the preferred lending program. Preferred lenders are screened by the SBA. The lenders apply to become a part of the program, and their regional SBA field office decides whether the lender has a competent performance record and is able to adequately judge borrower risk.[12] Lenders that are accepted into the preferred lender program are granted more autonomy in issuing loans. They must only submit an application to the SBA for an eligibility review but are otherwise responsible for the final credit decision. Thus, preferred lenders are more easily able to adjust their information acquisition practices.

For non-preferred lenders, the guarantee approval process consists of three broad steps. First, the lender submits an application to the SBA, which includes, for example, business information, expected use, current and projected financial information, and names of firm ownership.[13] The SBA reviews the application and then chooses whether to approve the guarantee. Conditional on approval, the remainder of the lending process is the same for preferred and non-preferred lenders. The decision moves to the borrower, which chooses whether to reject the offer or accept and pay the guarantee[14] and closing fees. Once the loan

[11] Congressional Research Service (2019a)

[12] SOP 50 10 5(E): Lender and Development Company Loan Programs, U.S. Small Business Administration Office of Finance

[13] https://www.sba.gov/loans-grants/see-what-sba-offers/sba-loan-programs/general-small-business-loans-7a/7a-loan-application-checklist%20

[14] Guarantee fees are technically the responsibility of the lender, but they typically pass this expense through to the borrower. For loans with maturity of 12 months or less, the guarantee fee is 0.25% of the guaranteed portion of the loan. For maturities over 12 months, the fee varies by loan amount. For loans of $150,000 or less, the fee is 2.0%. The fee is 3.0% for loans between $150,001 and $700,000, 3.5% for loans

is approved and disbursed, the borrower either repays each period or defaults. In the case of default, the SBA pays the lender the guaranteed percentage of the remaining loan balance.

It is important to note one further feature of the SBA 7(a) lending environment. Banks' pricing practices are constrained by the 7(a) program's interest rate caps, which vary by loan size and loan term (maturity). For example, during the period of analysis in this paper, a lender issuing a loan of more than $50,000 with a term of less than seven years may charge a maximum spread on a fixed-rate loan of 2.25% above the Fixed Base Rate.[15] Loans of the same size but with a maturity of more than seven years may be priced up to 2.75% above the Fixed Base Rate.[16] Given the presence of this constraint, banks may instead choose other contract characteristics, such as maturity, which can be up to ten years for non-real-estate loans and up to 25 years for real-estate loans, to maximize payoffs. In the remainder of the paper, I rely on a loan price that accounts for both interest rate and loan term (maturity). Computational details can be found in Appendix A.3.

Lender incentives are governed by the guarantee rate, and the SBA sets the maximum rates nationally, allowing them to vary only by loan amount. The standard bounds are 85% for loans of $150,000 or less and 75% for loans above $150,000. While guarantee rates are allowed to vary, and sometimes do so in practice, the majority are issued at the maximum level. In the empirical analysis, I exploit an exogenous increase in this maximum guarantee percentage during the recession. To stimulate lending to small businesses, Congress included guarantee increases in two pieces of stimulus legislation.[17]

---

between $700,001 and $1,000,000, and 3.75% for loans between $1,000,001 and $5,000,000. SOP 50 10 5(E): Lender and Development Company Loan Programs, U.S. Small Business Administration Office of Finance.

[15]The Fixed Base Rate is periodically published by the SBA. Prior to October 1, 2009, the Fixed Base Rate was equal to the prime rate. On October 1, 2009, the SBA changed the target rate from the prime rate to a function of the one-month LIBOR and the average of 5- and 10-year LIBOR swap rates. For details, see https://www.sba.gov/sites/default/files/bank_5000-1128.pdf. This policy change occurred outside the guarantee-change windows examined in this paper and is therefore not within the scope of the project.

[16]For a discussion of the price caps in the SBA 7(a) program across maturities and loan amounts, see https://www.sba.gov/sites/default/files/articles/SBA_Loan_Fact_Sheet_0.pdf.

[17]The changes to the guarantee program were part of a large legislative agenda. To ease concerns about confounding variation, such as changes to the SBA size standards leveraged in Denes et al. (2021), I show

The American Recovery and Reinvestment Act of 2009 temporarily increased the maximum rate to 90% for all loans.[18] This expansion began on March 16, 2009 and lasted until funding ran out on May 31, 2010. Later, the Small Business Jobs Act of 2010 again increased the maximum to 90% on September 27, 2010. In this case, the funding ran out on January 3, 2011.[19] In addition, in all expansion periods, guarantee fees were waived for all new loans. The Small Business Jobs Act also increased the maximum loan amount from $2 million to $5 million. These legislative changes provide four plausibly exogenous shocks (including both takeup and expiration) to government guarantees, which I leverage in my empirical analysis.

1.5 Descriptive Evidence

The guarantee-rate changes described above provide banks with incentives to lend to riskier borrowers at lower rates. Yet, the extent to which loan characteristics change depends on the elasticities described in Section 1.3 as well as any endogenous adjustments to the lenders' precision of information. In this section, I leverage the variation induced by the American Recovery and Reinvestment Act and Small Business Jobs Act to quantify changes to lending practices (i.e., loan characteristics, borrower characteristics, and prices) induced by higher guarantee rates. I find that lenders respond strongly to the adjustment in incentives. Banks bunch on the high-guarantee side of the policy changes and offer observably more generous loans (i.e., longer maturity, larger amount) at lower rates. They also appear to price risk less precisely, as the correlation between price and default, the ex-post outcome of interest, declines in the high-guarantee period.

that the empirical results are robust to examining only the expiration of the loan guarantees. The exact date at which the guarantees expired was a function only of the total amount of funding and was unrelated to any other programs that were instituted through the two legislative acts of interest.

[18] Because these changes aimed to stimulate new lending, the SBA did not allow loans approved prior to the guarantee increases to be cancelled and reapproved to receive a higher rate.

[19] Congressional Research Service (2019b)

1.5.1 Data

For the empirical analysis, I rely on a number of data sources. First, I collect publicly available loan data from the SBA 7(a) Loan Data Reports. This dataset contains loan-level information for all loans guaranteed through the 7(a) program and approved on or after January 1, 1990. From this source, I obtain loan characteristics (e.g., interest rate, term, amount borrowed, percent guaranteed), borrower and lender characteristics, and repayment outcomes. The dataset also includes cancelled applications which were approved for a guarantee by the SBA but cancelled prior to the first disbursement of funds. Cancellations occur automatically if the guarantee fee is not paid to the SBA within 90 days, but guarantees may also be cancelled if borrowers receive other funding options or close the business.[20] I use these cancelled applications as a measure of borrower rejection of a loan offer. For loans that were not cancelled, I assign a loan to the default classification if it is listed as charged off as of December 31, 2019. As of this date, only 18% of loans remain outstanding, and I group these loans with those that have been paid in full.

I augment the SBA 7(a) Loan Data Reports with lender information from the Federal Financial Institutions Examination Council (FFIEC) Bank Call Reports. This dataset contains bank balance sheet information and is released at a quarterly frequency. Because I use balance sheet information to capture variation across banks, rather than variation within bank across time, I consider data from a single point in time: December 31, 2010. Specifically, I use these data in the structural estimation to allow for heterogeneity in banks' costs of information acquisition and costs of lending. I match these characteristics to the loan data using the bank name and successfully match approximately 80% of loans.[21] The unmatched

[20]Dilger (2016) notes that the number of loan guarantees approved annually is higher than the number of loans disbursed. This report attributes these cancellations to changes in demand for funds, changes in business ownership, or changes in outside funding options. See https://nationalaglawcenter.org/wp-content/uploads/assets/crs/R41146.pdf.

[21]This match rate is approximately the same as that in Choi and Lee (2019). Note that this restriction yields a smaller sample for the structural portion (N = 11,447) than for the descriptive evidence (N = 13,992).

loans are drawn disproportionately from non-FDIC-insured institutions. While these lenders could face different incentives than their FDIC-insured counterparts, they are responsible for a small subset, a maximum of 20%, of SBA loans. Furthermore, the main takeaways from the descriptive evidence, using the full dataset, are identical to those in the structural analysis, which restricts to the matched loans. This alleviates selection concerns.

To complete the dataset, I obtain zip code level demographic data from the U.S. Census Bureau, Bureau of Labor Statistics, and Federal Housing Finance Agency to act as proxies for borrower characteristics. As with the FFIEC data, I rely on the demographic data to capture cross-sectional variation across geographic areas. For this reason, I obtain demographics from a single point in time: 2010. I merge information on local housing prices, population, and median household income.

For the main analyses, I restrict attention to loans of up to \$2 million[22] approved within a window of the four exogenous changes to guarantee rates and trim the price distribution at the 1st and 99th percentile to remove outliers.[23] Also, if multiple loans were issued to a single borrower between 2009 and 2011, I keep only the first loan. Because these analyses rely on an event-study approach using only temporal variation, focusing on a small window around events limits confounding factors, like changes to bank lending costs and other macroeconomic fluctuations affecting entrepreneurial activity, that could affect loan characteristics or outcomes.[24] The main specifications examine six-week windows around the exogenous changes, but the results are robust to a two-week window definition (see Appendix A.1). Additionally, I provide support in Appendix A.2 that changes to guarantee rates are not correlated with variables that determine bank lending costs and treasury yields, though

[22]The Small Business Jobs Act of 2010 increased the maximum loan amount from \$2 million to \$5 million. I exclude these newly eligible loans to standardize the borrower pool over time.

[23]In Appendix A.1, I show that the event-study results are robust to including the price outliers.

[24]For example, the prime rate remained constant at 3.25% from December 16, 2008 through December 17, 2015, covering the entire period of interest. See https://www.jpmorganchase.com/corporate/About-JPMC/historical-prime-rate.htm.

they are associated with changes to bank stock prices for a subset of the largest banks in the United States at the time. One final restriction relates to a specific feature of the second uptake of the guarantee expansion. A number of loans, whose applications were submitted in the months leading up to the guarantee expansion, were held in a queue to wait for higher guarantee rates. Because applications for these loans may have been submitted outside the small window around the guarantee expansion, I exclude them from the analysis.[25] Summary statistics for the final loan-level dataset are displayed in Table 1.1.

	Mean	S.D.	Min.	25th Pct.	Median	75th Pct.	Max.
All Loans							
Interest Rate (Pct.)	5.85	0.56	2.75	5.5	6	6	9.03
Loan Return (Yearly, Pct.)	2.13	0.39	1.31	1.75	2.23	2.38	3.29
Term (Months)	164.33	88.50	7	90	120	244	318
Amount Borrowed (Thousands)	558.23	488.09	6.5	200	400	773.05	2,000
Guaranteed Pct.	0.86	0.06	0.32	0.85	0.9	0.9	0.9
Acceptance	0.87						
Loan Size > 150,000	0.81						
Preferred Lender	0.71						
Observations	13,992						
Accepted Loans							
Default	0.07						
Observations	12,159						

Table 1.1: Summary Statistics

Note: This table displays summary statistics for the main sample of loans issued within six weeks of guarantee-rate changes. It relies solely on data from the SBA 7(a) Loan Data Reports.

1.5.2 Lending Activity Over Time

The SBA data paint a clear picture of the lenders' response to the guarantee-rate increases. Banks adjust the extensive margin of lending, and an expansion of lending activity is ap-

[25] The SBA cleared the entire queue within one week of the expansion on September 27, 2010. For this reason, I exclude loans issued within seven days of that event. See https://www.sba.gov/about-sba/sba-newsroom/press-releases-media-advisories/sba-loan-queue-cleared-one-week-after-signing-jobs-act.

parent in time series plots. Figure 1.3 examines the trend in lending over time. Panel (a) displays total approvals and Panel (b) displays the share of loans issued by preferred lenders. Together, they illustrate the stark response in lending behavior to guarantee-rate changes. More generous guarantees result in a large spike in lending activity, and higher average amounts of lending persist throughout the period. Furthermore, the bunching around events is correlated with a higher share of preferred lenders. These banks are provided more autonomy in the lending process and are therefore more able to shift loans into the high guarantee period.

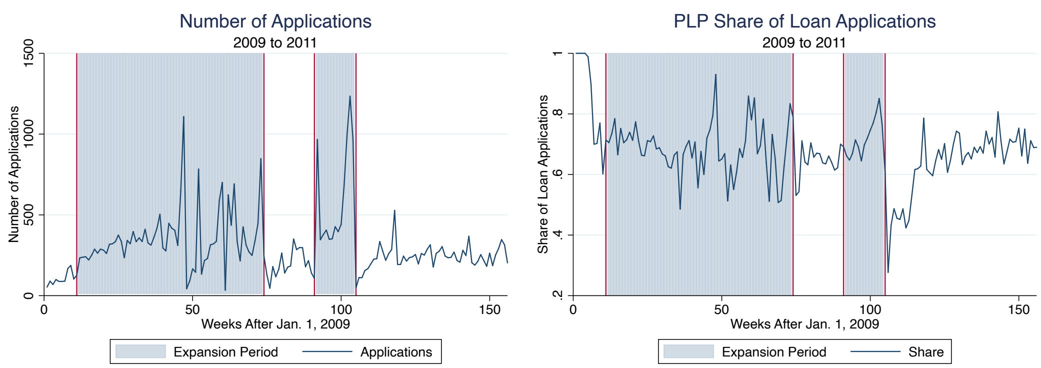

(a) Applications (b) PLP Share

Figure 1.3: Lending Activity Over Time

Note: This set of figures displays time series plots, from 2009 to 2011, for (1) the number of total 7(a) guarantee applications and (2) the share of loan applications filed by institutions in the Preferred Lender Program (PLP). The analysis relies on data from the SBA 7(a) Loan Data Reports.

Taken together, these trends suggest that loans are not necessarily assigned randomly into treatment, and, for this reason, the event-study results presented in the remainder of this section should not be interpreted as causal effects. Instead, I intend for them to describe the equilibrium responses to the policy change, combining the guarantee pass-through, the information effect, and any shift in the distribution of borrowers. The results, and their shortcomings, help motivate the structural analysis, which allows me to disentangle these three pieces of the observed equilibrium changes.

1.5.3 Event Studies

Not only do lenders issue a larger number of loans to capture higher government guarantees; they also offer contracts with more generous terms. To quantify the extent to which lenders alter loan characteristics in the high-guarantee periods, I estimate event-study specifications which exploit the temporal policy variation. The regressions all take the following form:

$$Y_{ijt} = \alpha + \delta \mathbb{I}(t = \text{SBA Recovery}) + \beta X_{it} + \epsilon_{ijt},$$

where Y_{ijt} is the relevant loan characteristic (i.e., interest rate, yearly return, amount borrowed, loan term) for loan i issued by lender j in period t, and X_{it} is a matrix of borrower covariates. These covariates include business type, NAICS code (two-digit), event date fixed effects (i.e., corresponding to each introduction and expiration of the guarantee expansion), and an indicator of whether the loan is used for real estate (i.e., maturity $>$ 10 years), as well as zip code demographics, including the change in the housing price index since 2008, median household income, and total population. All demographic variables are normalized to be between 0 and 1. δ is the coefficient of interest and captures the changes to average loan characteristics during the guarantee expansion period.

Table 1.2 displays event-study results for loan characteristics. Interest rates and yearly loan return (a measure of the NPV of all loan payments, assuming full amortization)[26] decline in the guarantee expansion period, while the average amount borrowed and loan term increase. All of these results are both statistically and economically significant. The average interest rate decreases by approximately 4.2 basis points, the average loan amount increases by \$56 thousand, and the average loan term increases by between 4 and 5 months, all of which are consistent with the relaxation of credit constraints on small businesses. Conditional on observables, lenders are willing to provide financing at lower prices, whether

[26] See Appendix A.3 for a full description of the yearly return calculation. I assume full amortization of the amount borrowed and discount payments at a rate equivalent to 5% per year.

it be through lower interest rates, larger loans, or longer terms, when government guarantees are higher.

Furthermore, these results are driven by the response of preferred lenders. Since preferred lenders are afforded almost complete discretion in the lending process, it is natural to think that they are better equipped to respond to policy variation by changing their information acquisition practices and loan characteristics. In Appendix A.4, I show that the equilibrium response to the guarantee expansion is large and statistically significant for preferred lenders, while the response for non-preferred lenders is small and typically statistically insignificant.[27] In the structural analysis, I allow for lender-side heterogeneity by preferred status to capture any differences in information acquisition costs stemming from the program rules.

The lower interest rates, longer maturities, and larger capital outlays suggest that a portion of the gains induced by the guarantee program are passed on to borrowers (i.e., the guarantee pass-through is weakly positive), on average. However, as described in Section 1.3, this pass-through may differ across the risk distribution. Moreover, the extent to which borrowers benefit also depends on the magnitude of the information effect, which controls how precisely individual loan characteristics match the borrower's underlying risk. If all borrowers receive more generous loan outcomes, they are unambiguously better off with loans in the 7(a) program and, in turn, are less likely to accept their outside option. Alternatively, if information effects increase the likelihood that low-risk borrowers receive offers in the upper tail of the price distribution, then these borrowers are disproportionately harmed by the policy change and could be more likely to opt out of the guarantee program and instead take their outside option. Without a more explicit model of the lending market, I am unable to

[27] To assuage concerns that the differential response could be driven by observable differences between the two groups, I provide evidence in Appendix A.4 that preferred lenders are not geographically concentrated. I do find suggestive evidence that these lenders are larger, as measured by the value of all loans on their balance sheets. Because of this, I provide one further robustness check and restrict the sample to include only banks in the top quartile of the distribution of loan value. Using this sample, the results are qualitatively equivalent, as I again find larger responses to the policy change by preferred lenders. This eases concerns that the results are driven by bank characteristics rather than participation in the Preferred Lender Program.

	(1) Interest Rate (Pct.)	(2) Loan Return (Yearly, Pct.)	(3) Amt. Borrowed ($ Thousands)	(4) Loan Size > 150,000	(5) Loan Term (Months)
Loans Issued Within 42 Days of Events					
SBA Recovery	-0.0420**	-0.0329***	56.40***	0.0572***	4.241***
	(0.0163)	(0.00700)	(10.96)	(0.00879)	(0.945)
Mean Outcome	5.85	2.13	558.23	0.81	164.33
Observations	13,992	13,992	13,992	13,992	13,992
Zip Code Dem. Controls	✓	✓	✓	✓	✓
Business Type FE	✓	✓	✓	✓	✓
NAICS (Two-Digit) FE	✓	✓	✓	✓	✓
Real Estate FE	✓	✓	✓	✓	✓
Event Date FE	✓	✓	✓	✓	✓

Standard errors are clustered by lender.
* p<0.1, ** p<0.05, *** p<0.01

Table 1.2: Event-Study Results – Loan Characteristics

Note: This table presents results of the main event-study specification, including all loans issued within 42 days of guarantee-rate changes. All specifications include controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date.

directly measure changes to price across the distribution of borrower types. Before turning to the model, I provide descriptive evidence supporting the notion that lenders price risk less precisely (i.e., there is more mispricing of risk), and high-risk borrowers receive larger price changes than their safer counterparts, when loan guarantees decrease the cost of default faced by the lender.

1.5.4 Loan Pricing

The event-study results are informative of the impact of increases in guarantee rates on equilibrium loan characteristics, but they do not allow for analysis of heterogeneous effects across the risk distribution and do not speak to the information channel directly. To examine these two points, I exploit two features of the SBA lending setting: (1) lender pricing decisions are informative of their information structure and belief about a borrower's risk and (2) observed ex-post repayment outcomes are informative of true borrower risk.[28] Changes in the mapping from borrower default risk to loan price, above what can be explained by product characteristics, are suggestive of differences in the lender's precision of information and can be used to test whether the observed guarantee policy variation disproportionately impacts certain types of borrowers.[29] I examine changes in this mapping during the high-guarantee period and provide evidence that lenders price risk less precisely, which is consistent with a decline in their screening effort and the disproportionate benefits accruing to high-risk borrowers.

This section proceeds in two steps. First, I provide evidence that the correlation of price and default declines in the high-guarantee period. In the second piece of the analysis, I illustrate the heterogeneity in price effects across the borrower risk distribution. In particular, I show that low-risk borrowers receive higher price offers in the high-guarantee period than they do in the baseline, while the opposite is true for high-risk borrowers. These two results motivate the need to consider the distributional impact of guarantee policies in addition to average outcomes, as aggregate effects fail to capture the heterogeneous impact of both the guarantee pass-through and any changes in lenders' information precision.

Price-Default Correlation

To examine how guarantee-rate changes affect the correlation of price and default, I rely on a two-stage regression framework. In the first stage, I estimate a flexible mapping from

[28]These features are frequently present in markets with asymmetric information and are exploited in related ways in Chiappori and Salanie (2000), Rajan et al. (2015), and Crawford et al. (2018).

[29]It is important to note that a change in the mapping from default to price could also stem simply from heterogeneity in the guarantee pass-through. As described earlier, the value to the lender of insurance against ex-post default risk varies depending on the borrower's underlying risk. A higher guarantee is less valuable to a lender extending a loan to a low-risk borrower than it is to a lender that offers funding to a high-risk business. Thus, a change to the guarantee rate should affect prices differently across the distribution of borrower risk. The analysis presented in Appendix A.4 provides evidence that the observed changes in price-default mappings are not the result of only this mechanical effect. If changes in the precision of information played no role, we would expect the mappings of preferred and non-preferred lenders to move together. The fact that preferred lenders, who are more able to adjust their information acquisition, respond disproportionately suggests the presence of an information effect.

loan characteristics to prices using observations in both the low- and high-guarantee periods. Specifically,

$$p_{ijt} = f(M_{ijt}, B_{ijt}) + X_{ijt}\beta + \epsilon_{ijt},$$

where p_{ijt} is the loan price offered to borrower i by lender j in period t. To standardize the price across loans, I compute the price as the NPV equivalent for a ten-year loan, as described in Appendix A.3. This same price is used in the structural analysis. f is a flexible function (i.e., quadratic terms and interactions) of M_{ijt}, the guarantee rate provided by the SBA and B_{ijt}, the size of the loan. X_{ijt} is a vector that includes fixed effects for borrower characteristics, such as business type and two-digit NAICS code, as well as event-date fixed effects and a control for loans with maturity over 10 years. I also control for zip code demographics using the same covariates as in the event-study specifications. ϵ_{ijt} captures all borrower, loan, and lender characteristics priced into the loan but not observed by the econometrician. This residual could include soft information collected by the loan officer, hard information not included in the SBA dataset, or lender characteristics, including its marginal cost of lending.

The second stage estimates the correlation between the unobserved component of price, ϵ_{ijt}, and borrower default, the observed ex-post loan outcome. I estimate the following specification:

$$d_{ijt} = \gamma_1\epsilon_{ijt} + \gamma_2\mathbb{I}(t = SBA) + \gamma_3\epsilon_{ijt} \times \mathbb{I}(t = SBA) + g(M_{ijt}, B_{ijt}) + X_{ijt}\delta + e_{ijt},$$

where d_{ijt} is an indicator of default, $g(\cdot)$ is a flexible function of the guarantee rate and loan size (again, quadratic terms and interactions), and X_{ijt} is the same vector of controls as in the first stage. The coefficient of interest is γ_3, which measures the change in the correlation between ex-post default and the unobservable component of price during the SBA Recovery period. In essence, this coefficient captures any changes in the lenders' precision of risk pricing (or, potentially, screening effort). A negative coefficient indicates that loan prices are less informative of borrower risk in the high-guarantee period than in the low-guarantee period.

	(1) Default	(2) Charge Off
ϵ	2.210	1.367
	[1.762,2.688]	[1.059,1.653]
$\mathbb{I}(t = SBA)$	-0.027	-0.025
	[-0.055,0.003]	[-0.043,-0.005]
$\epsilon \times \mathbb{I}(t = SBA)$	-0.653	-0.470
	[-1.172,-0.212]	[-0.776,-0.172]
Raw Correlation	0.203	0.181
SD(ϵ)	0.030	0.030
Observations	12,159	12,159

Table 1.3: Pricing Regression Results

Note: This table presents results for the second stage of the two-stage pricing regression. The first column displays results for an indicator of default on the lefthand side, while the second column displays results for the share of the loan charged off. All specifications include quadratic terms and interactions of the guarantee rate and loan amount and controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date. Block-bootstrapped (by lender) 95% confidence intervals are displayed in brackets; N=1,000.

Table 1.3 displays estimates for the second stage. In aggregate, there is a positive correlation between the unobserved components of price and default. Intuitively, lenders observe more than the econometrician and therefore are better able to price borrower risk. The correlation falls substantially during the high-guarantee period, with the coefficient on the soft-information term falling by almost 30% in the SBA Recovery period, indicating a change in the mapping from borrower risk to price. As in the event studies, the response of preferred lenders differs from that of their counterparts. Appendix A.4 presents the bootstrap distribution of γ_3 for both types of lenders. As expected, preferred lenders are better able to adjust their pricing function, owing to the greater autonomy provided them by the SBA.

Heterogeneity Across the Risk Distribution

The equilibrium decline in the correlation between price and default, captured by the results in Table 1.3, proves that the mapping from risk to price changes in response to policy variation, but it does not address the question of which borrowers benefit and which, if any, are harmed. In this section, I examine in more detail price changes across the distribution of risk and show that high-risk borrowers receive substantially lower prices in the high-guarantee period than in the baseline. In contrast, low-risk borrowers receive slightly higher prices, but the change is statistically indistinguishable from zero.

Figure 1.4 displays the default rate, relative to the mean rate of the period, for each quartile of the residualized price distribution (i.e., ϵ) for loans in the low- and high-guarantee periods. This illustration approximates the mappings from borrower risk, as determined from ex-post outcomes, to prices in each period. From this plot, it is clear that the price effects are heterogeneous across the risk distribution. Low-risk borrowers receive slightly higher price offers in the high-guarantee period than they do in the baseline, though this result is noisy. The opposite is true for borrowers on the upper tail of the risk distribution, and this difference is statistically and economically significant. This finding implies that the guarantee pass-through to low-risk borrowers is either (1) negative (i.e., leads lenders to offer higher prices when guarantees increase in generosity) or (2) sufficiently weak to be dominated by changes in the precision of the lenders' signals. In the structural analysis, I quantify the relative magnitudes of each of these potentially competing effects.

Taken together, these descriptive results show that lenders respond swiftly to the policy changes. In the high-guarantee period, lenders are willing to issue observably riskier loans at cheaper prices. The average interest rate falls, the average loan size increases, and the average loan maturity increases. Lenders also appear to price risk less precisely, as the correlation between price offers and default declines. These event studies capture average effects of loan guarantees but are not fit for analysis of compositional changes and price-risk mismatch.

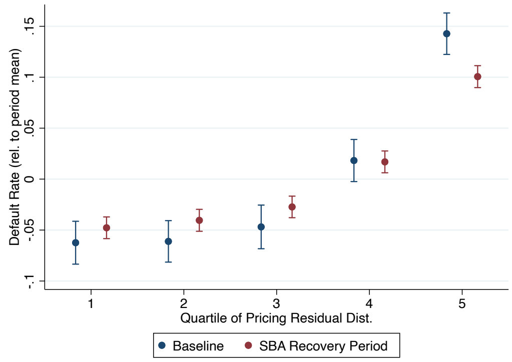

Figure 1.4: Price Effect Heterogeneity – By Borrower Risk

Note: This figure displays the mean default rate relative to the average rate in the specified period for loans in each quintile of the pricing residual distribution (i.e., the distribution of ϵ). The dots indicate point estimates, while the bars denote the 95% confidence interval.

A model allows for the estimation of primitives governing borrower and lender decisions, which in turn informs the analysis of borrower risk composition and pricing. Furthermore, with a model, I am able to separately quantify the guarantee pass-through and changes to lenders' precision of information. In the remainder of the paper, I present a model of the small-business lending market, describe the estimation procedure, and present counterfactual exercises using the estimates of the underlying primitives.

1.6 Model and Estimation

The empirical model retains the main elements of the simplified model presented in Section 1.3, but provides a general framework, in the style of Crawford et al. (2018), to allow for flexible correlation between borrower risk and willingness to pay as well as observable heterogeneity on both the borrower and the lender side. This heterogeneity captures, for example, differences in default and acceptance decisions across business types (e.g., by NAICS code) and locations. It also allows for pricing and information acquisition practices to differ according to whether a bank is a preferred lender.

In the remainder of this section, I describe each component of the model in detail. I then discuss identification of the primitives governing borrower acceptance and default decisions, as well as the lenders' information and pricing rules.

1.6.1 Framework

As before, let borrowers be indexed by i and lenders be indexed by j. Conditional on observables, borrowers are characterized by their propensity to repay, ξ_i^R, and their price responsiveness, α_i, distributed jointly by $F_{\xi^R,\alpha}$. The lender does not observe the borrower's repayment type, ξ_i^R, but chooses its signal precision and obtains a noisy signal, s_{ij}, of the borrower's risk. Lenders are heterogeneous in the cost they must pay to gain a unit of signal precision and thus choose different optimal information acquisition strategies. Conditional on the signal the lender receives, it chooses the loan characteristics to offer, and borrower outcomes (i.e., acceptance and default) are then realized. A more detailed description of the timing of the model is as follows:

1. Borrower i, whose type (risk) and price responsiveness are private information, is paired with lender j.

2. With knowledge of the guarantee rate and its cost of information acquisition, lender j chooses the precision of its information and receives a signal of borrower i's type. Conditional on the signal, lender j offers a loan of price p_{ij} to the borrower.

3. The borrower then receives a shock to its value of receiving a loan or, analogously, a shock to the value of the outside option, and borrower i chooses whether to accept the loan.

4. Conditional on acceptance, the borrower defaults or repays according to its propensity to repay, ξ_i^R.

One key assumption of the modeling strategy is the existence of lender market power, through the exogenous pairing of borrower i and lender j. Two institutional details support this assumption. To receive a loan through the SBA 7(a) program, a borrower must pass the credit elsewhere test, meaning lenders face limited, if any, competition from non-SBA lenders for 7(a) loans. SBA lenders may compete with one another, but I provide evidence that such competition is limited. I observe all loans approved by the SBA for a guarantee and find few instances in which a single borrower is approved for a guarantee with multiple lenders. In Appendix A.5, I conduct a number of analyses to provide further support for this assumption.[30]

Borrower Acceptance and Default

As in Crawford et al. (2018), I model the borrowers' acceptance and default decisions using reduced-form rules. In this framework, the baseline level of adverse selection is determined by the correlation of ξ_R and α, rather than by a structural relationship mapping borrower risk to the acceptance decision, as in Stiglitz and Weiss (1981). Suppose businesses have the following utility of loan repayment:

$$u_i^R = X_i^R \beta^R + \xi_i^R,$$

[30] A number of studies have examined the connection between competition and information acquisition, including Ruckes (2004), Hauswald and Marquez (2006), and Avramidis et al. (2021). These models show that information acquisition incentives and competition are negatively related (i.e., incentives are strongest when competition is low). The fact that I do not model competition would, therefore, be problematic if (1) competition played a role in information-acquisition decisions and (2) the level of competition changed with the generosity of guarantees. In this case, my model would attribute differences in information precision across guarantee schemes as a response to changes in the lenders' risk exposure rather than a response to changes in competition. As described in the main text, the institutional details of my setting and analyses shown in Appendix A.5 suggest that strategic considerations are limited and therefore this misattribution does not appear to be a concern.

where X_i^R is a vector of borrower covariates that shift the probability of repayment[31] and ξ_i^R is the borrower's private-information propensity to repay. I normalize the utility of default to zero, and borrowers repay if $u_i^R \geq 0$.

One assumption underlying the repayment decision merits further discussion. I impose that default rates are independent of prices, conditional on observables and ξ_i^R. This assumption is common in the literature on SBA-guaranteed lending (see, e.g., Cox et al. (2021)), as well as work on other forms of consumer credit (see, e.g., Nelson (2022)). Allowing for moral hazard would temper price effects and could introduce non-monotonicities into the mapping from signals to prices.[32]

The borrowers' acceptance decisions take a similar form, as the utility they receive from obtaining a loan is a borrower-specific function of covariates, price, and a preference shock:

$$u_{ij}^A = X_i^A \beta^A - \alpha_i p_{ij} + \epsilon_{ij},$$

where X_i^A is a vector of borrower covariates that shift the probability of acceptance, p_{ij} is the price offered by lender j to borrower i, and ϵ_{ij} is a preference shock, which I assume to be distributed Type I Extreme Value. Adverse selection enters this model through the correlation of ξ_i^R and α_i, where riskier borrowers may be less price-sensitive than their safer counterparts.[33]

[31] Given this structure, I implicitly assume that the level of collateral is captured by borrower observables, X_i^R, and factored into the mean return of the project. A change in the level of collateral would alter the default payoff and would be indistinguishable from a level shift in the utility of repayment. While some loans in this program are secured with collateral, the stated purpose of the program is to expand credit to parties that do not satisfy the requirements of a typical bank lending policy. One of these requirements is sufficient collateral, so the assumption that collateral levels are captured by observables is likely less binding here than in other lending contexts.

[32] Depending on the shape of the default function, non-monotonicities could arise if the degree to which expected default is responsive to changes in price is sufficiently different across signal realizations. For a further discussion of these non-monotonicities, see Crawford et al. (2018).

[33] In Appendix A.8, I show that results are robust to a model with a random coefficient on the constant instead of price.

Lender Information Acquisition and Pricing

The lenders' decision problem proceeds in two stages. First, lenders choose the precision of information to collect about each borrower. Suppose this amounts to choosing the joint distribution of signals and borrower risk, denoted F_{s,ξ^R}. Then, lenders receive signals and offer prices to maximize ex-ante expected profits, conditional on the signal. As each lender solves the problem using backward induction, I describe the stages in reverse order, beginning with the price offers.

At this stage, lenders take as given the precision of information and the joint distribution of borrower risk and price sensitivity, $F_{\xi^R,\alpha}$, in the population. Lender j then receives a noisy signal of borrower i's risk, ξ_i^R, and updates its beliefs using Bayes' rule. I assume the signal has the following structure:

$$s_{ij} = \xi_i^R + \sigma_{\gamma,ij}\epsilon_{\gamma,ij},$$

where $\sigma_{\gamma,ij}$ is the standard deviation of the signal noise for the loan from lender j to borrower i and $\epsilon_{\gamma,ij} \sim F_{\epsilon_\gamma}$ is a mean-zero, symmetrically distributed error.

The lender sets its price offer to maximize its ex-ante expected payoff, conditional on the signal. Specifically, the lender solves:

$$\max_{p_{ij}} \int\int P^A(\alpha_i, p_{ij}) \left[(1-(1-M_{ij})P^D(\xi_i^R))p_{ij} - \zeta_{ij}\right] \, dF_{\xi_i^R,\alpha_i|s_{ij}},$$

where $P^A(\alpha_i, p_{ij})$ is the probability of acceptance for a borrower with price sensitivity α_i at a price of p_{ij}, $P^D(\xi_i^R)$ is the probability of default for a borrower of risk ξ_i^R, ζ_{ij} is the marginal cost of lending, and M_{ij} is the guarantee rate. Rearranging the first-order condition of the lender's pricing problem, the optimal price is characterized by the following differential equation:

$$p^*_{ij} = \underbrace{\zeta_{ij} \frac{\int\int \frac{\partial P^A}{\partial p}(\alpha_i, p^*_{ij})\, dF_{\xi^R_i,\alpha_i|s_{ij}}}{\int\int \frac{\partial P^A}{\partial p}(\alpha_i, p^*_{ij})(1-(1-M_{ij})P^D(\xi^R_i))\, dF_{\xi^R_i,\alpha_i|s_{ij}}}}_{\text{Guarantee-corrected effective MC}} - \underbrace{\frac{\int\int P^A(\alpha_i, p^*_{ij})(1-(1-M_{ij})P^D(\xi^R_i))\, dF_{\xi^R_i,\alpha_i|s_{ij}}}{\int\int \frac{\partial P^A}{\partial p}(\alpha_i, p^*_{ij})(1-(1-M_{ij})P^D(\xi^R_i))\, dF_{\xi^R_i,\alpha_i|s_{ij}}}}_{\text{Guarantee-corrected captured markup}}$$

The resulting pricing function captures key features of the risk-based pricing framework set forth by Phillips (2013) and Edelberg (2006), as prices are a function of lending costs and a markup term, and lenders charge a premium depending on the borrower's risk. The pricing function also captures the two effects of a guarantee-rate change described in Section 1.3. The guarantee pass-through manifests itself through a change in the lender's payoff under default, which shifts both the effective marginal cost and the expected capture of the markup. The information effect enters through the joint distribution of risk and price responsiveness, conditional on the signal. If information is more precise, prices are more responsive to changes in the realized signal.

To close the model, I specify the process by which lenders obtain information. With knowledge of the above pricing rule, lenders choose the precision of their information ex ante. Lenders observe borrower characteristics, e.g., location and NAICS code, and are informed of the joint distribution of borrower risk and price sensitivity. These observables shift the lender's expected payoff as well as the marginal benefit of additional information. The lenders then choose how much costly effort they will expend to collect more information about the borrower. This extra effort could be used to obtain "soft" information, such as a subjective measure of trustworthiness, through face-to-face meetings with the borrower or to verify other codifiable information like revenue projections and business plans.

In terms of the model, the lender's choice amounts to setting the precision of the noisy signal, or $\sigma_{\gamma,ij}$. I assume that precision is costly to acquire and that lenders must pay $C_{ij}(\sigma_{\gamma,ij}) = \kappa_{ij} \cdot \frac{1}{\sigma^2_{\gamma,ij}}$ to obtain a signal with standard deviation $\sigma_{\gamma,ij}$. If lenders must meet with borrowers to acquire information, we would expect the information cost to scale linearly with the number of meetings and the number of meetings to positively correlate with the final precision. Pomatto et al. (2020) provide an axiomatic foundation for information acquisition that microfounds this form of cost structure.[34]

To choose its signal precision, a lender maximizes its expected profit, given borrower observables and the guarantee rate. The optimal precision solves

$$\max_{\sigma_{\gamma,ij}} \int \int \int P^A(\alpha_i, p_{ij})[(1-(1-M_{ij})P^D(\xi_i^R))p_{ij} - \zeta_{ij}]\, dF_{\xi_i^R,\alpha_i|s_{ij};\sigma_{\gamma,ij}}\, dF_{s_{ij};\sigma_{\gamma_{ij}}} - C_{ij}(\sigma_{\gamma,ij}),$$

where, as before, $P^A(\cdot,\cdot)$ is the probability of loan acceptance, $P^D(\cdot)$ is the probability of default, ζ_{ij} is the cost of lending, and $C_{ij}(\cdot)$ is the cost of information acquisition. The precision of information affects payoffs through two components. First, it enters into the probability of default, as the lender updates its beliefs of borrower risk. With information that is more precise, prices and borrower types are more highly correlated. Second, it affects lender beliefs over acceptance decisions, through the correlation of price responsiveness and borrower risk.

The equilibrium level of information equates the marginal benefit associated with more precise risk pricing and the cost of acquiring a signal of that precision. As shown in the stylized model, these benefits vary with the guarantee rate, the borrower's risk, and how strongly acceptance decisions respond to price changes. To examine the responsiveness of information to guarantee-rate changes and the extent to which the subsequent price-risk mismatch yields a differential impact on certain types of borrowers, we must estimate the

[34]Whereas the standard mutual information-based costs in rational inattention models capture the costs of processing information, the function used here instead captures the costs of acquiring information. The axioms set forth by Pomatto et al. (2020) ensure monotonicity of costs in precision and imply that the marginal cost of information is constant. In my setting, these are both attractive properties.

primitives of the model. In the remainder of this section, I describe how I take the model to the data.

1.6.2 Identification and Estimation

Parameterization

I now provide details on the parameterization of the model and discuss identification of the primitives. I specify the borrower's utility of repayment as a linear function of borrower observables, X_i^R. This matrix includes a constant, the normalized loan amount[35], indicators for loans with maturity less than or equal to ten years, event-date fixed effects, two-digit NAICS fixed effects, business type fixed effects (i.e., partnership, individual, or corporation), and normalized zip code level demographics, such as the change in the housing price index, median household income, and total population.

The borrower's utility of acceptance is a linear function of a subset of the above borrower observables. Namely, this utility varies by the normalized loan amount, whether maturity is less than or equal to ten years, and event date. Also, I allow the utility of acceptance to change in the SBA Recovery period. This accounts for the fact that guarantee fees, which are typically passed through to the borrower, were waived whenever guarantees were expanded.

To complete the borrower side, I parameterize the joint distribution of risk and price responsiveness as a joint normal. This parameterization requires two normalizations. I set $\mathrm{E}[\xi_i^R] = 0$ and $\mathrm{Var}[\xi_i^R] = 1$, so the repayment decision takes the form of a standard probit. The remaining parameters governing the joint distribution are to be estimated. Specifically, the distribution takes the following form:

$$\begin{pmatrix} \xi_i^R \\ \alpha_i \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ \mu_\alpha \end{pmatrix}, \begin{pmatrix} 1 & \rho\sigma_\alpha \\ \rho\sigma_\alpha & \sigma_\alpha^2 \end{pmatrix}\right)$$

[35]All normalized variables, w_i, are computed as $\frac{w_i - \min_i w_i}{\max_i w_i - \min_i w_i}$.

On the lender side, I assume the marginal cost function is linear and separable in the shifters of lending cost, Z_{ij}, the ongoing guarantee fee, which is a function of the guarantee rate, and an independent shock, ω_{ij}. Specifically,

$$\zeta_{ij} = \beta^Z Z_{ij} + \psi(M_{ij}) + \omega_{ij},$$

where the shifters of lending cost, Z_{ij}, include normalized bank-level interest-bearing balances and non-interest-bearing balances, currency, and coin due from depository institutions, as well as bank-state level balances. This allows for cost variation across banks, as well as within banks, across states. Banks with larger balances have a lower cost of acquiring funds and, therefore, lower prices. The ongoing fee is 0.55% per year for the guaranteed portion of the loan, and I provide details of the computation of $\psi(M_{ij})$ in Appendix A.3. I assume $\omega_{ij} \sim N(0, \sigma_\omega^2)$, where σ_ω is a parameter to be estimated.

Finally, as mentioned above, I assume that the lender updates its beliefs of the borrower's type using Bayes' rule. The lender receives a signal of borrower risk:

$$s_{ij} = \xi_i^R + \sigma_\gamma(h_j)\epsilon_{\gamma,ij},$$

where the standard deviation of the noise of the signal, $\sigma_\gamma(h_j) = \exp(\beta^{\sigma_\gamma} h_j)$, and h_j, a vector of lender covariates, consists of a constant, an indicator for preferred lenders, an indicator for the SBA Recovery period, and an interaction term. This parameterization allows lenders to respond to guarantee-rate changes through the precision of their information and also accounts for the fact that preferred lenders respond differentially, indicating a difference in information acquisition costs. I assume $\epsilon_{\gamma,ij} \sim N(0, 1)$.

Under this parameterization, the lender computes the joint posterior distribution of ξ_i^R

and α_i, conditional on s_{ij}, as

$$\begin{pmatrix} \xi_i^R \\ \alpha_i \end{pmatrix} \sim N\left(\begin{pmatrix} \tilde{\mu}_{\xi^R} \\ \tilde{\mu}_\alpha \end{pmatrix}, \begin{pmatrix} \tilde{\sigma}^2_{\xi^R} & \tilde{\rho}\tilde{\sigma}_\alpha \\ \tilde{\rho}\tilde{\sigma}_\alpha & \tilde{\sigma}^2_\alpha \end{pmatrix}\right),$$

where

$$\tilde{\mu}_{\xi^R} = \frac{s_{ij}}{1+\sigma_\gamma(h_j)^2},\ \tilde{\sigma}^2_{\xi^R} = \frac{\sigma_\gamma(h_j)^2}{1+\sigma_\gamma(h_j)^2},$$

$$\tilde{\mu}_\alpha = \mu_\alpha + \frac{s_{ij}\rho\sigma_\alpha}{1+\sigma_\gamma(h_j)^2},\ \tilde{\sigma}^2_\alpha = \frac{\sigma^2_\alpha(\sigma_\gamma(h_j)^2 + (1-\rho^2))}{1+\sigma_\gamma(h_j)^2}, \text{ and}$$

$$\tilde{\rho} = \rho\frac{\sigma_\gamma(h_j)^2}{1+\sigma_\gamma(h_j)^2}.$$

Identification

I now discuss identification of the model primitives. We observe borrower acceptance and default decisions, lender pricing decisions, and observables that vary across both borrowers (e.g., NAICS code and geography) and lenders (e.g., bank balances). The primitives of interest are components of the borrowers' utility of repayment, the components of their utility of acceptance, and the joint distribution of borrower types, price responsiveness, and signals.

On the borrower side, parameters are informed by variation in both default and acceptance rates. Under the location and scale normalizations described above, variation in default rates across the covariates entering the borrowers' utility of repayment pins down the coefficients of the repayment utility function. Similar variation pins down the coefficients on the exogenous shifters of the utility of acceptance (i.e., those entering X_i^A).

Because prices are set with knowledge of s_{ij}, which is not independent of ξ_i^R and α_i, another source of variation is required to recover the joint distribution of risk and price responsiveness, $F_{\xi^R,\alpha}$. I rely on cost shifters, in the form of the balance sheet instruments described in Section 1.6.2, to provide the requisite variation. The mean of the distribution of price responsiveness is informed by differences in acceptance rates across banks with different balance sheet sizes and, within banks, across states with different levels of branch balances. The variance of this distribution is pinned down by the extent to which price responsiveness varies across borrowers that differ in their exogenous covariates. Finally, ρ is informed by the correlation between acceptance decisions and residualized default. For example, given a set of exogenous covariates, if borrowers that unexpectedly accept the loan (i.e., low α_i) are also much more likely to default (i.e., low ξ_i^R), then ρ is positive and large in magnitude. I impose one assumption on ρ, which I detail below.

Assumption 1: $\rho \geq 0$.

This assumption restricts the direction of selection in the market. Namely, it ensures that demand is not advantageously selected. Under advantageous selection, prices need not be monotonic in the signals received by lenders. Lenders learn about price responsiveness only through its correlation with risk, ξ_i^R. If advantageous selection is sufficiently strong, then a borrower that is believed to be high risk may receive a lower price offer because its level of risk implies weak demand. A similar, non-advantageous selection assumption is imposed in the consumer lending literature, e.g., Nelson (2022).

The lender-side parameters are informed by pricing decisions, and, in particular, variation in prices across exogenous borrower covariates, lender cost-side instruments, and the borrowers' residual risk. Under the assumption that ω_{ij} is an independent, i.i.d. cost shock, the linear parameters of the marginal-cost function are identified by variation in average price across the cost shifters, Z_{ij}. The remaining parameters ($\sigma_\gamma(h_{ij})$ and σ_ω) are informed by two sources of variation: (1) variation in prices across borrower covariates and default status and (2) the variance of prices.

The intuition for the separate identification of the standard deviation of the distribution of signal noise, $\sigma_\gamma(h_j)$, and that of the i.i.d. cost shock, σ_ω, stems from differential pricing across ex-post default outcomes. A similar intuition underlies the analysis in Section 1.5. In this exercise, suppose we condition on X_i^R and X_i^A, the borrower covariates that determine repayment and acceptance, Z_{ij}, the exogenous cost shifters, h_j, the lender covariates entering the standard deviation of the signal noise, and M_{ij}, the guarantee rate. In the remainder of this section, I suppress these conditioning variables for notational convenience.

Define $a_{ij} = 1$ if borrower i accepts the loan from lender j and zero otherwise. Let d_{ij} take the same values for the default decision. Consider two conditional distributions: $p_{ij}|a_{ij} = 1, d_{ij} = 1$ and $p_{ij}|a_{ij} = 1, d_{ij} = 0$. The first is the distribution of price offers for borrowers that default in the end, and the second is the distribution of price offers for borrowers that do not. Note that prices are set at loan origination, while the default outcome is realized ex post. Thus, a given price offer is informative of what lender j believes about borrower i's risk, while the ex-post default outcome is a proxy for that borrower's risk.

To set ideas, first assume there is no variation in lending costs. As prices are decreasing in the signal realization, a decline in the standard deviation of the signal noise results in an increase in the separation of the two conditional distributions. Intuitively, prices respond more to borrower risk, proxied with the ex-post default outcome, when information is more precise. It follows that the precision of information is informed by the difference in the location of these two distributions. This simple intuition relies on the lack of variation in lending costs, as part of the difference between the location of the distribution conditional on default and that conditional on repayment could be due to cost shocks. To disentangle variation in marginal costs from the lenders' precision of information, I rely on the facts that (1) I observe not only the locations of the two distributions but also their spread, and (2) the cost shocks are independent of borrower risk. It follows that changes in the variance of cost shocks correspond to changes in the width of both conditional distributions (e.g., an increase in the variance of the cost shocks increases the widths of both the distribution of prices conditional on default and that conditional on repayment). Together, the location and spread of the distributions of price offers across ex-post outcomes pins down the information precision and the variance of lending costs. With the important sources of variation in hand, I next describe the estimation procedure.

Estimation

For each borrower, I observe the triplet (a_{ij}, d_{ij}, p_{ij}), where a_{ij} denotes the acceptance decision, d_{ij} denotes the default decision, and p_{ij} is the observed price offer. As described in Appendix A.3, I compute the price as the NPV of a normalized ten-year loan. Let Θ denote the full vector of parameters to be estimated: $\Theta = [\beta^R, \beta^A, \beta^Z, \sigma_\omega, \mu_\alpha, \sigma_\alpha, \rho, \sigma_\gamma(h_{ij})]$.

I estimate the model using full-information maximum likelihood, integrating over the joint distribution of types, price responsiveness, and signals. More details on the estimation procedure can be found in Appendix A.6.

1.7 Estimates of Borrower and Lender Primitives

Using the maximum likelihood procedure, I recover estimates of the primitives governing both borrower and lender decisions. This section consists of three main components. First, I present estimates of the borrower primitives and discuss how they capture the reduced-form responses illustrated by the event-study analyses, as well as other variation in default and acceptance decisions across borrower types. Second, I discuss estimates of lending costs and the lenders' information quality. Lenders respond to higher guarantees by decreasing the precision of information they collect about borrower risk. In the high-guarantee period, the average signal-to-noise ratio declines by 8.3%, though this effect varies by lender type (i.e., preferred vs. non-preferred). Third, I exploit the first-order condition of the lender's ex-ante profit function to back out the implied cost of information acquisition. These primitives guide the counterfactual responses I describe in Section 1.8.

1.7.1 Estimates

Borrower-Side Results

Borrower characteristics capture substantial variation in acceptance and default both within and across periods. Figure 1.5 displays estimates of individual-level acceptance and default probabilities, based solely on observed borrower covariates, in the baseline and SBA Recovery period. The majority of borrowers accept loans between 80 and 95 percent of the time, and the rates decline after the policy change. This variation across periods suggests a difference in borrowers' outside options in the SBA Recovery period. Default probabilities range from close to zero to approximately 20 percent. The borrower covariates again capture variation in repayment across periods, as observably safer borrowers comprise a larger share of loans after the policy change than in the baseline.

Within period, borrower observables, including the primary industry of their business (i.e., two-digit NAICS code), the loan amount, and the time at which they request funding (i.e., the event date), contribute to heterogeneity in acceptance and default rates. These results reflect differences across the observed covariates in the typical use of the loaned funds, or in the level of demand for a particular good or service, among other differences. Below, I describe in more detail the sources of variability in the borrowers' acceptance and default decisions.

Table 1.4 lists selected parameter estimates and standard errors for borrower-side primitives. The full set of estimates is presented in Appendix A.7. The borrowers' utility of repayment has a noisy relationship with the loan amount. Businesses likely complete different types of risky projects depending on the size of the loan, but it appears that there is still

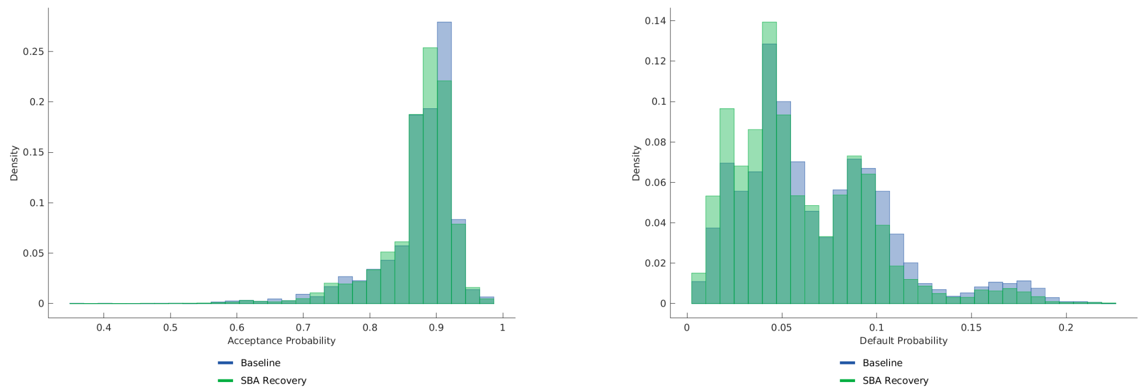

Acceptance Probability Default Probability

Figure 1.5: Estimates of Acceptance and Default Probabilities by Period

Note: These figures display the predicted probability of default, conditional on X_i^R, and the predicted probability of acceptance, conditional on X_i^A. The distributions are induced by the sample distributions of X_i^R and X_i^A rather than by variance in the predicted probabilities.

significant variation in the projects, conditional on the loan amount. One other driver of heterogeneity in default is the expected use of the loan, which I proxy using discrete categories of loan maturity. Real estate loans (i.e., those with maturity longer than ten years) default at approximately 44% the rate of loans with maturities less than ten years. There is also substantial variation in the utility of repayment across time. Specifically, loans issued at the height of the recession in late 2009 have an average default rate of 10.7%, while those issued in the midst of the recovery in 2010 and early 2011 default at rates between 5 and 6%. Not only does mean return vary across loan characteristics and time; it also differs across borrower characteristics. Loans issued to businesses in the Finance & Insurance sector (NAICS Code 52) and the Real Estate sector (NAICS Code 53) default at lower rates – 1.4% and 1.7% on average, respectively – than firms in other industries, which default at an average rate of 6.4%.

Borrowers also face different outside options across time and across guarantee schemes. Borrowers that receive larger loans have lower utility of acceptance, or better outside options.

This result may seem counterintuitive given the existence of the credit elsewhere test, as it should be easier to acquire a small amount of capital from a non-bank source. However, borrowers that request large loans may have easier access to personal or family funds than those that seek smaller loans. Additionally, they may have access to other funding sources, such as venture capital. There is also pronounced variation in the value of the outside option depending on whether the loan is used for real estate. Borrowers accept real estate loans at a rate of 86.5%, while they accept shorter, non-real estate loans at a rate of 87.9%. Given the availability of collateral for real estate loans, borrowers may find it easier to obtain financing from sources outside the 7(a) program. Finally, there is noisy evidence of temporal variation in the value of the outside option, as borrowers that receive loans in 2010 and 2011 accept at lower rates than those that receive loans in 2009. This difference across time captures the procyclical nature of the availability of outside funds.

The final parameters of interest govern the joint distribution of borrower risk and price responsiveness. The mean and standard deviation of the marginal distribution of α imply an average residual elasticity of -2.72 across all borrowers. Furthermore, price responsiveness and propensity to repay are positively correlated, though the estimate of the correlation coefficient is noisy. Taken together, these borrower-side results imply that there is substantial heterogeneity, both observable and unobservable, in borrowers' acceptance and default decisions. This variation across borrowers informs the lenders' pricing decisions and precision of information, and I present the results of the lender-side estimation in the next section.

Lender-Side Results

Lenders' decisions are governed by their marginal cost function, the associated structural pricing rule, and their choices of information quality. Figure 1.6 displays the distribution of marginal costs, including a simulated cost error, ω_{ij}. The estimates of the primitives underlying the cost function can be found in Appendix A.7. It is important to note that

Parameter	Estimate (S.E.)
Components of β^R:	
Constant	1.434***
	(0.215)
Normalized Amt. Borrowed	-0.084
	(0.095)
Maturity $\leq$ 10 Years	-0.419***
	(0.049)
Event 2	0.396***
	(0.061)
Event 3	0.302***
	(0.060)
Event 4	0.373***
	(0.055)
NAICS: 52 (Finance & Insurance)	0.618*
	(0.347)
NAICS: 53 (Real Estate & Rental & Leasing)	0.386
	(0.277)
Components of β^A:	
Constant	22.310***
	(6.647)
Normalized Amt. Borrowed	-0.479***
	(0.168)
Maturity $\leq$ 10 Years	1.309***
	(0.455)
Event 2	-0.018
	(0.017)
Event 3	-0.021
	(0.017)
Event 4	-0.050**
	(0.023)
Components of $F_{\xi^R,\alpha}$	
μ_α	16.682***
	(4.987)
σ_α	0.701
	(1.032)
ρ	0.655
	(0.697)

* p<0.1, ** p<0.05, *** p<0.01

Table 1.4: Selected Estimates – Borrower-Side Parameters

Note: This table presents estimates of selected borrower-side parameters. Standard errors, calculated as described in Appendix A.6, are displayed in parentheses. Full results can be found in Appendix A.7.

the cost estimates imply that banks are lending one dollar at a cost of less than one dollar. There are a number of potential explanations for this finding, and I discuss a few of the most plausible in this section.

First, lenders may anticipate future profits from the relationship when extending funds through the SBA 7(a) program. There are switching costs associated with changes in banking relationships. If the nascent businesses that receive loans through the 7(a) program survive their early years, they may rely on the same bank for deposit accounts and other loans in the future. This provides incentives for banks to lend at a discount at the outset, ensuring they capture the relationship and any ensuing profits. Second, lenders may face reputational costs for opting out of the 7(a) program. Thus, they choose to participate, despite price caps limiting the markups they are able to charge.

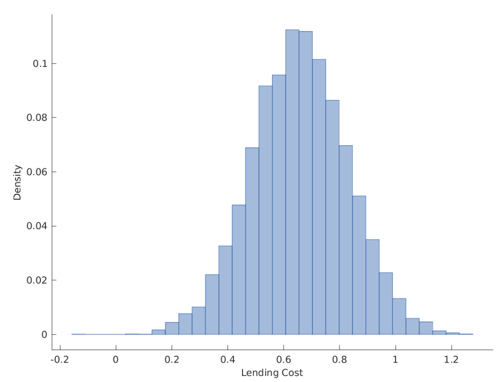

Figure 1.6: Estimates of Lending Cost, ζ_{ij}

Note: The costs displayed in this figure include the shock, ω_{ij}. The distribution captures variation across banks in cost shifters, Z_{ij}, and the magnitude of the shock. It does not capture variance in the estimates of β^Z.

The second aspect of the lender's decision problem is the choice of information precision. As mentioned in Section 1.6, I allow the signal distribution to vary across periods (baseline vs. SBA Recovery) and types of lenders (preferred vs. non-preferred). The heterogeneity in information structures capture: (1) any endogenous changes to signal precision in response to an increase in the guarantee rate and (2) differences in levels of information and responsiveness to guarantees across borrower types. The latter is crucial in this setting, as preferred lenders must meet SBA lending standards to be admitted to the program (i.e., meet a threshold information acquisition cost or signal precision) and are given substantially more autonomy in the lending process, providing them more scope to adjust their information acquisition practices in response to changes in incentives.

Table 1.5 presents estimates of the standard deviations of the lenders' signal distributions, by both period and lender type. The average lender chooses a noisier signal structure in response to higher guarantee rates. This result suggests the presence of an indirect information effect, and the magnitude of the aggregate effect is sizable. Across all banks, the standard deviation of the signal distribution moves from 1.311 in the baseline to 1.402 in the SBA Recovery period, on average. This corresponds to a decrease in the signal-to-noise ratio $\left(\frac{1}{1+\sigma_\gamma^2}\right)$ from 0.368 to 0.337 (-8.3%). In the following subsection, I provide more context about this effect and examine the pass-through of information-precision changes to prices offered to borrowers from different quantiles of the risk distribution.

It is also important to note that there is considerable heterogeneity across lender types in information-acquisition practices, owing to the observable differences in the response to the policy changes between preferred and non-preferred lenders (see, e.g., Section 1.5). In the baseline and after the policy change, non-preferred lenders obtain more precise signals of borrower quality than their counterparts in the preferred program. However, preferred and non-preferred lenders not only obtain different levels of information; they also respond differently to guarantee-rate changes. The estimates of σ_γ imply that the signal-to-noise ratio decreases in the SBA Recovery period by 11.5% for preferred lenders, while it increases by 8.4% for non-screened lenders, though the latter effect is noisy and statistically indistinguishable from zero. This differential response captures the fact that, before being allowed to close a loan, preferred lenders require only a simple review by the SBA to ensure that a loan is eligible for a guarantee (e.g., the borrower passes the credit elsewhere test and requires

Parameter	Estimate (S.E.)
S.D. of Signal Distribution: σ_γ	
Non-Preferred Lender, Baseline	1.230***
	(0.108)
Preferred Lender, Baseline	1.355***
	(0.104)
Non-Preferred Lender, SBA Recovery	1.148***
	(0.109)
Preferred Lender, SBA Recovery	1.485***
	(0.093)
Difference Across Periods (SBA Recovery - Baseline):	
Non-Preferred Lender	-0.082
	(0.055)
Preferred Lender	0.130***
	(0.048)

* p<0.1, ** p<0.05, *** p<0.01

Table 1.5: Selected Estimates – Lender Information

Note: This table presents estimates of the parameters governing the precision of the lenders' information. Standard errors, calculated as described in Appendix A.6, are displayed in parentheses. Full results can be found in Appendix A.7.

funding for an allowed purpose under the SBA 7(a) program). On the other hand, non-preferred lenders must complete a standardized application packet to ensure that the due diligence process adheres to SBA requirements. While these banks may still have some room to alter their screening practices, the lack of complete autonomy constrains, and potentially even eliminates, their ability to shirk.

The estimates of signal precision map to unique costs of information (i.e., the unit cost of additional signal precision, κ_{ij}), which I back out using the first-order conditions of the lenders' ex-ante expected profit maximization problems. At the time of the decision, the lender has not received a signal of the borrower's propensity to repay and thus takes the expectation of profits over the joint distribution of signals, borrower types, and price responsiveness. I allow the cost to vary at the loan level. Let $\tau_{ij} = \frac{1}{\sigma^2_{\gamma,ij}}$. I recover costs by solving the following first-order condition:

$$\frac{\partial}{\partial \tau_{ij}} \int \int \int P^A(\alpha_i, p_{ij})[(1-(1-M_{ij})P^D(\xi_i^R))p_{ij} - \zeta_{ij}]\, dF_{\xi_i^R,\alpha_i|s_{ij};\sigma_{\gamma,ij}}\, dF_{s_{ij};\sigma_{\gamma_{ij}}} = \frac{\partial C}{\partial \tau_{ij}} = \kappa_{ij}.$$

To operationalize this process, I perturb $\sigma_{\gamma,ij}$ around its estimated value and compute a finite-difference estimate of the marginal revenue term on the lefthand side, which is known given the estimates of borrower and lender primitives. Figure 1.7 displays the loan-level estimates of the cost of information acquisition.

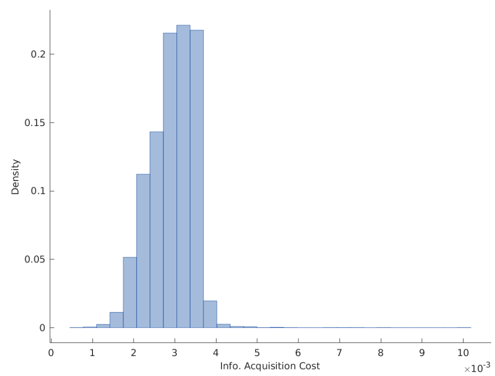

Figure 1.7: Estimates of Information Cost, κ_{ij}

Note: The distribution displayed in this figure captures variation in borrower and lender covariates (i.e., X_i^R, X_i^A, Z_{ij}, and ω_{ij}). These observables map to the marginal profit of information precision and, therefore, imply different marginal costs of information, κ_{ij}.

For non-preferred lenders in the SBA Recovery Period, the average marginal cost of information corresponds to an outlay of 17 basis points for a normalized ten-year loan. For the average loan of \$558,230, the lender pays \$928 to acquire information about that borrower. The total costs are lower for preferred lenders. For the same average loan amount, these lenders pay \$805 to obtain a signal of the borrower's risk.

1.7.2 Model Fit

Before analyzing the predictions of the model in counterfactuals, I demonstrate that the model fits the observed outcomes well. In this section, I pay particular attention to the extent to which the model captures variation in acceptance, default, and pricing decisions across policy regimes. Properly estimating how responsive decisions, and particularly the choice of price offers, are to guarantee-rate changes is essential both for disentangling the guarantee pass-through and information effect and for simulating counterfactual guarantee policies. Table 1.6 displays observed and simulated moments by policy regime. The model captures changes across periods in average prices and average default rates and predicts the correct level of acceptance in aggregate, though the response to the policy change is a bit muted. The model not only fits aggregate moments but also captures variation at the loan level in price offers. Figure 1.8 displays overlaid distributions of observed and simulated prices.

	Baseline		SBA Recovery	
Moment	Obs.	Sim.	Obs.	Sim.
$E[p_{ij}]$	1.242	1.240	1.233	1.232
$E[a_{ij}]$	0.904	0.882	0.865	0.881
$E[d_{ij} \| a_{ij} = 1]$	0.084	0.066	0.060	0.064
$E[a_{ij}p_{ij}]$	1.122	1.094	1.067	1.086
$E[d_{ij}p_{ij} \| a_{ij} = 1]$	0.107	0.085	0.076	0.081

Table 1.6: Model Fit

Note: This table displays estimates of relevant moments using the observed and simulated data for both the Baseline and SBA Recovery periods. I denote the price offer as p_{ij}, an indicator of acceptance as a_{ij}, and an indicator of default as d_{ij}.

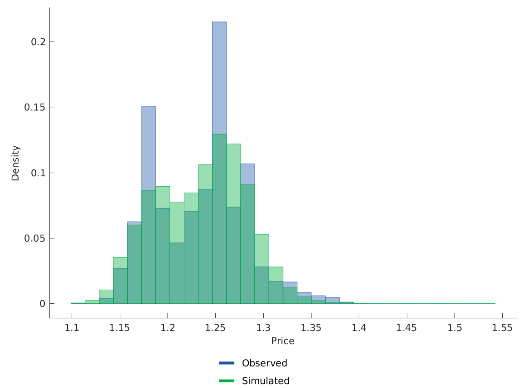

Figure 1.8: Observed vs. Simulated Prices

1.8 Guarantees, Information, and Policy Design

The results of the previous sections illustrate the existence of a guarantee pass-through as prices decline, on average, in the high-guarantee period, as well as an indirect effect that manifests itself through the information channel. The relative magnitudes of each of these components determines the equilibrium impact of an increase in the guarantee rate on prices. Specifically, for low-risk borrowers, the information effect leads lenders to offer higher prices and works in the opposite direction of the program's intended effect.

In this section, I quantify the extent to which the guarantee pass-through and information effect contribute to equilibrium changes in price and borrower surplus in response to variation in the guarantee rate. I find that, in response to an increase in the generosity of guarantees, the information effect amplifies the price decline (and borrower-surplus increase) for high-risk borrowers (i.e, those in the bottom quintile of the distribution of u_i^R). However, lenders exploit their market power on low-risk borrowers and raise prices in response to more generous guarantees. The information effect exacerbates these price increases and the corresponding borrower-surplus declines.

Given the fact that lenders respond through both channels and that such behavior benefits high-risk borrowers at the expense of their counterparts, it is natural to ask whether an alternative policy leads to similar price declines without disproportionately benefitting borrowers at the top of the risk distribution. A hybrid policy, with a less generous guarantee and a cost subsidy, could do so through two channels. First, such a policy leaves lenders with more ex-post risk exposure, leading them to collect information that is more precise. Second, the subsidy decreases lenders' marginal costs, inducing lower prices across the distribution of borrowers. I find that, holding expected government spending fixed, a hybrid policy generates gains in aggregate borrower surplus. Under a policy with a 50% guarantee and a cost subsidy such that expected spending remains fixed, high-risk borrowers are slightly worse off than under the baseline of a 90% guarantee and no subsidy, but the gains to low-risk borrowers more than offset the losses. In aggregate, borrower surplus increases by between 1.5 and 2%. This result suggests that there exist policy mechanisms that limit moral hazard effects and, consequently, temper the distributional impact of guarantee programs.

1.8.1 Decomposition of Price and Borrower-Surplus Effects

To disentangle the guarantee pass-through from the information effect, I perform a decomposition exercise. I restrict attention to loans issued in the SBA Recovery period. This homogenizes the sample in terms of observables and eliminates any changes to observable borrower composition in the recovery period. For guarantee rates $\tilde{M} \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1\}$, lenders solve their ex-ante profit maximization problem given their marginal cost of information acquisition, κ_{ij}, their marginal cost of lending ζ_{ij}, setting the NPV of the ongoing guarantee fee to be constant at its baseline value $(\tilde{M} = 0.9)$[36], and the primitives governing the joint distribution of borrower risk and price responsiveness. Let $\sigma^*_\gamma(\tilde{M})$ denote the optimal standard deviation of the lender's signal noise distribution for a guarantee rate of $\tilde{M}$.

[36]Fixing the ongoing fee throughout this exercise isolates the impact of the guarantee pass-through and information effect. If I were to allow the fee to vary with the guarantee rate, changes in lending cost would confound the magnitude of the guarantee pass-through.

The decomposition exercise then consists of two main portions. First, I simulate (ξ_i^R,α_i) for each borrower from the estimated distribution (i.e., using the estimates of μ_α, σ_α, and ρ). Then, I draw signals and compute prices and borrower surplus for each $\tilde{M} \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1\}$ under two different scenarios:

1. Standard deviation of signal noise fixed at $\sigma_\gamma^*(0.9)$

2. Standard deviation of signal noise set at its optimal level, $\sigma_\gamma^*(\tilde{M})$

For the remainder of this section, I refer to the outcomes for $\tilde{M} = 0.9$ as the baseline. To quantify the magnitude of the pass-through of guarantees, I examine the changes in price across guarantees under Scenario 1 above. These changes hold the level of information fixed and thus isolate the direct effect of guarantees. The information effect is the difference between the price computed under Scenario 2 and that computed under Scenario 1 and captures the additional impact of bank moral hazard on price offers.

In each stage of the decomposition, I compute a measure of borrower surplus using the standard log-sum formula (see, e.g., Train (2009)), and scale by α_i so its units are equivalent to those for prices:

$$BS_{ij} = \frac{1}{\alpha_i} \log\left(1 + \exp(X_i^A \beta^A - \alpha_i p_{ij})\right).$$

Then, to scale these values to dollars over the normalized ten-year loan term, I multiply this value by the loan amount, b_{ij}.

For all levels of guarantee generosity, $\tilde{M}$, lenders respond to guarantee-rate changes by adjusting the level of information they collect. The average level of information obtained, in equilibrium, is monotonically decreasing in $\tilde{M}$ for $\tilde{M} \in \{0.5, 0.6, 0.7, 0.8, 0.9, 1\}$. Figure 1.9 displays the average signal-to-noise ratio across guarantee rates. Under a guarantee rate of 50%, the signal-to-noise ratio is, on average, 22% higher than under the baseline of 90%. An increase in the generosity of the guarantee decreases the marginal value of information. Thus, under the assumption that marginal acquisition costs are fixed, the declining relationship between the signal-to-noise ratio and the guarantee rate is expected.

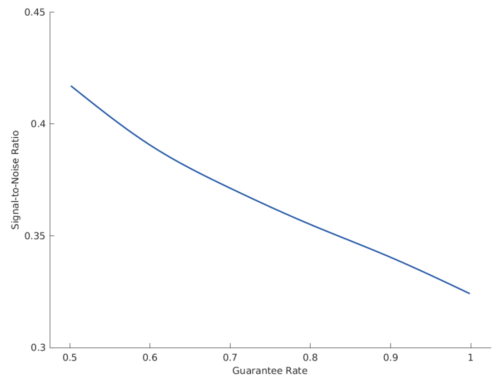

Figure 1.9: Average Signal-to-Noise Ratio vs. Guarantee Rate

Note: This figure displays, for the specified guarantee rates, the average signal-to-noise ratio, across all loans.

This change in information structure interacts with the guarantee pass-through to determine equilibrium outcomes. Figure 1.10 displays the equilibrium effects on price and borrower surplus of changes to the guarantee rate and quantifies the contribution of the guarantee pass-through and information effect to price changes, where the baseline is taken to be the optimal price under a guarantee rate of 90%. The top panel displays the average price and average borrower surplus, in the NPV of dollars over 10 years, across guarantees. The bottom panel decomposes the guarantee pass-through and information effect and displays, at the borrower level, the average percent change in price attributed to each of these two components.

This figure illustrates two key results. First, for the average borrower, price decreases and borrower surplus increases with the generosity of guarantees. When lenders are allowed to flexibly adjust their precision of information, the average price under a guarantee rate of 50% is 0.2% higher (and borrower surplus 0.7% lower) than that under the baseline guarantee of 90%. Second, the guarantee pass-through and information effect work in the

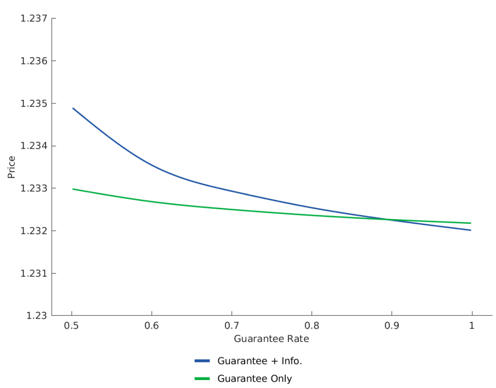

(a) Average Prices vs. Guarantee Rate

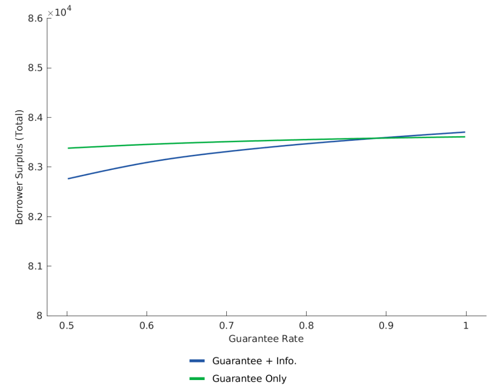

(b) Borrower Surplus vs. Guarantee Rate

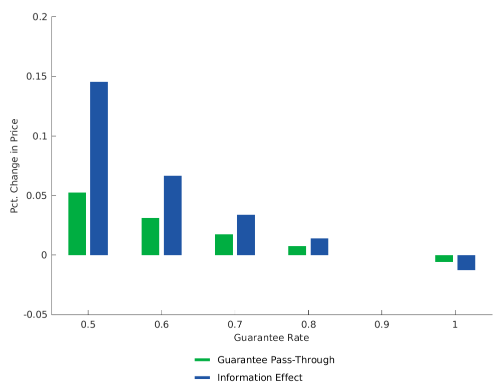

(c) Guarantee P-T, Info. Effect on Prices

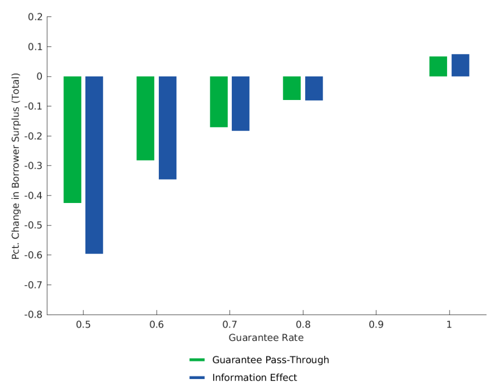

(d) Guarantee P-T, Info. Effect on Borrower Surplus

Figure 1.10: Decomposition of Price and Borrower Surplus Changes

Note: The top panel displays average loan price and borrower surplus across guarantee rates. In the bottom panel, the guarantee pass-through measures the percent change in price or borrower surplus under a change to only the guarantee rate (i.e., holding the precision of information constant). The information effect measures the additional change when information endogenously adjusts. The estimates displayed in the bottom panel are the average percent changes, again computed at the borrower level.

same direction for the average borrower, as the information effect magnifies the impact of guarantee-rate changes on price offers. Holding the lenders' signal precision constant, prices under a guarantee rate of 50% would be less than 0.1% higher (and borrower surplus approximately 0.4% lower), on average, than under the baseline of 90%. The adjustment of information precision amplifies these effects, contributing an additional 0.1% increase in prices and 0.6% decline in borrower surplus, on average.

However, as discussed, these average effects mask heterogeneity across the distribution of borrower risk. High-risk borrowers benefit disproportionately from generous guarantees, and this could come at the expense of their safer peers. Figure 1.11 displays the decomposition exercise on prices (top panel) and borrower surplus (bottom panel) separately for borrowers in the top quintile of the utility of repayment and those in the bottom quintile. High-risk borrowers, in the bottom quintile, benefit from more generous guarantees, while those in the top quintile (low-risk borrowers) are harmed. Furthermore, the guarantee pass-through and information effect magnify one another for both sets of borrowers. Holding the lenders' information structure constant, prices are approximately 0.3% higher, and borrower surplus 1.9% lower, for high-risk borrowers (i.e., those in the bottom quintile of the distribution of utility of repayment, u_i^R) under a guarantee of 50% than under the baseline of 90%. For low-risk borrowers (i.e., those in the top quintile), the less generous guarantee of 50% yields prices that are less than 0.1% lower and borrower surplus that is 0.3% higher. Allowing for flexible information acquisition amplifies both of these effects. In total, high-risk borrowers face 1.1% higher prices under a guarantee of 50% than under the baseline, while low-risk borrowers receive 0.2% lower price offers, on average. These price differences correspond to sizable changes in borrower surplus, as high-risk borrowers' surplus declines by 6.6%, on average, and that of low-risk borrowers increases by 2.5%.

These results highlight stark heterogeneity in the incidence of loan guarantee programs. The risk protection, by itself, leads to gains for high-risk borrowers and losses for safer borrowers, and bank moral hazard amplifies both effects. While guarantee schemes are popular in small-business lending markets, they are not the only instrument available to policymakers. Generous guarantees benefit the average borrower, but not all borrowers benefit. This result suggests there is room for alternative policy designs that may temper the effects of bank moral hazard and offset a portion of the losses absorbed by low-risk borrowers. In the next subsection, I examine one such policy, a combination of a guarantee

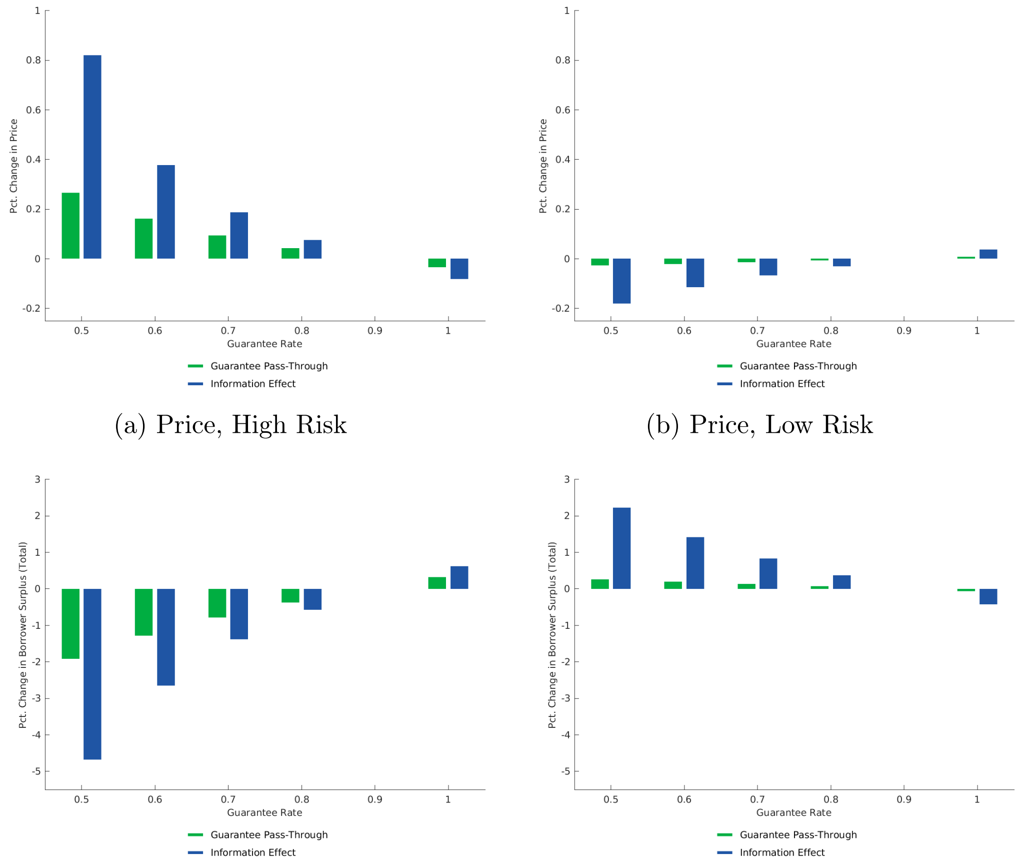

(a) Price, High Risk (b) Price, Low Risk

(c) Borrower Surplus, High Risk (d) Borrower Surplus, Low Risk

Figure 1.11: Decomposition of Price and Borrower Surplus Changes, By Borrower Risk

Note: The guarantee pass-through measures the percent change in price or borrower surplus under a change to only the guarantee rate (i.e., holding the precision of information constant). The information effect measures the additional change when information endogenously adjusts. The estimates displayed are the average percent changes, computed at the borrower level.

and a subsidy.

1.8.2 Hybrid Policy Design: Subsidy and Guarantee

The results of the decomposition exercise show that both the guarantee pass-through and the information effect benefit high-risk borrowers at the expense of their safer counterparts. This motivates examining ways to moderate the distributional impact of the program by providing incentives that act against the moral hazard problem. Given the observed difference in signal precision across lender types (preferred vs. non-preferred), one way to do this would be to increase oversight and remove the preferred lenders' autonomy. However, preferred lenders exist for a reason, and it may be prohibitively costly for the SBA to conduct an extensive review of each loan application.

Another alternative is to provide an additional lender incentive to counteract the heterogeneous price response due to moral hazard. A cost subsidy increases expected profits for all borrower types. This places downward pressure on prices across the distribution of risk and could therefore alleviate some of the losses borne by low-risk borrowers. In this section, I show that combining such a subsidy with a less generous guarantee, which limits the informational response, leads to gains over the baseline policy with a guarantee of 90%.

I examine a set of policies in which lenders are provided a subsidy and a guarantee below the baseline of 90% with the subsidy set such that expected government outlays are equal to those under the baseline guarantee and no subsidy. Specifically, define the subsidy, $S(\tilde{M})$ to solve:

$$\sum_{ij} b_{ij} P^A(p_{ij}(0.9), \alpha_i) \left[0.9 \cdot P^D(\xi_i^R) \cdot p_{ij}(0.9)\right] = \sum_{ij} b_{ij} P^A(p_{ij}(\tilde{M}), \alpha_i) \left[\tilde{M} \cdot P^D(\xi_i^R) \cdot p_{ij}(\tilde{M}) + S(\tilde{M})\right],$$

where $p_{ij}(\tilde{M})$ is the price offered by lender j to borrower i under a guarantee of $\tilde{M}$, and marginal costs are therefore given by: $\zeta_{ij} = \beta^Z Z_{ij} + \psi(0.9) + \omega_{ij} - S(\tilde{M})$.

As guarantees become less generous and subsidies increase, lenders collect more precise information about borrowers than they do in the baseline. However, the magnitude of the change in signal-to-noise ratio is lower than it would be absent the subsidy. Figure 1.12 displays the change in signal-to-noise ratio for each alternative policy (guarantee + subsidy), as well as the change under only an adjustment of the guarantee (i.e., similar to the exercise considered in Section 1.8.1), when compared to the baseline guarantee of 90%. Under a policy with a guarantee rate of 50% and a subsidy, lenders' average signal-to-noise ratio is approximately 21% higher than under the status quo, which is slightly lower than the 23% increase under only a guarantee-rate adjustment. While this change in signal precision is weaker under the hybrid policy, in the remainder of this section, I show that the pass-through of the subsidy offsets, in aggregate, any losses from the noisier information.

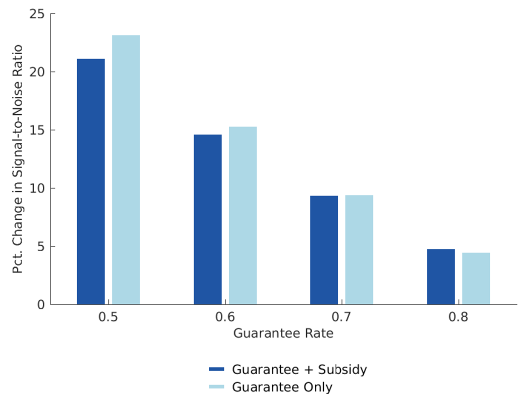

Figure 1.12: Percent Change in Signal-to-Noise Ratio for Hybrid Policies

Note: This figure displays the average change in the signal-to-noise ratio for (1) a hybrid policy with the stated guarantee rate and a subsidy set such that expected spending is the same as in the baseline (i.e., 90% guarantee and no subsidy) and (2) a policy with the stated guarantee rate and no subsidy.

The higher precision of information leads to gains for low-risk borrowers and losses for their higher-risk peers, dampening the distributional impact of the guarantee program. Figure 1.13 displays changes in prices and borrower surplus relative to the baseline guarantee of 90% for (1) a hybrid policy with a guarantee of 50% and a subsidy defined as described above and (2) a guarantee-only policy with a rate of 50%. The comparison of price offers and borrower surplus under these schemes concisely captures the forces to consider when designing guarantee policies. I report extensions of these counterfactual results, considering alternative combinations of guarantees and subsidies, in Appendix A.9. The qualitative takeaways are identical.

The plots in Figure 1.13 display the average percent change across all borrowers and also display changes for borrowers in the top and bottom quintiles of the distribution of u_i^R. A hybrid policy with a guarantee rate of 50% and a subsidy set such that expected spending does not change leads to lower prices and higher borrower surplus, on average. Prices decline by approximately 0.1% relative to the status quo, while borrower surplus increases by 1.6%. However, not all borrowers are better off. High-risk borrowers experience a decline in borrower surplus of close to 4%, though this decline is smaller than they would have faced under a guarantee-only policy. The aggregate gains accrue to lower-risk borrowers. Borrowers in the top quintile of the distribution of u_i^R experience gains in borrower surplus of about 5.5%, and this offsets a portion of the heterogeneous impact of the status quo program. The resulting shift toward a less risky borrower composition decreases the aggregate default rate by 0.1 percentage point (1.6%).

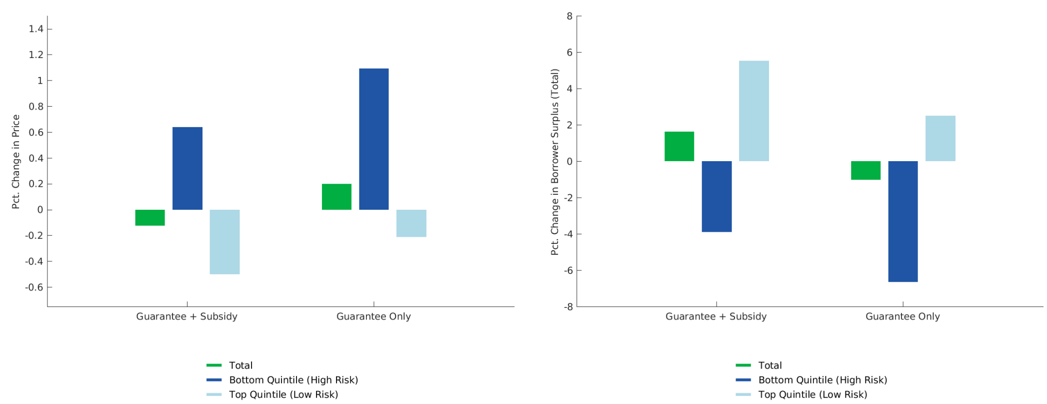

(a) Price (b) Borrower Surplus

Figure 1.13: Changes to Price and Borrower Surplus Under Hybrid Policies

Note: This figure displays the average change in price and borrower surplus for (1) a hybrid policy with a guarantee rate of 50% and a subsidy set such that expected spending is the same as in the baseline (i.e., 90% guarantee and no subsidy) and (2) a policy with a 50% guarantee rate and no subsidy. The estimates are displayed separately for all borrowers, high-risk borrowers, and low-risk borrowers.

It is important to note that this hybrid policy result does not indicate an improvement in social welfare. There are a number of reasons why policymakers may place disproportionate weight on the surplus of high-risk borrowers (e.g., these borrowers are more innovative or they tend to be individuals from demographic groups that traditionally found it difficult to obtain credit). But, the exercise demonstrates an important tradeoff in lending markets. Policymakers can induce lenders to extend funds at lower prices through two main channels: by providing ex-post insurance against default and by directly subsidizing lending costs. In the presence of costly information acquisition, these two methods can differ in their distributional impact. I show that a hybrid policy that includes both a subsidy and a guarantee outperforms the status quo policy of only a guarantee when judged on an aggregate borrower-surplus standard. Moreover, these surplus gains are the result of a decline in the heterogeneity of the program's impact across the distribution of borrower risk.

1.9 Conclusion

Policymakers frequently intervene in small-business lending markets in an attempt to expand financing to typically credit-constrained borrowers. Loan guarantee programs are widespread, both in the U.S. and elsewhere, and governments frequently rely on them to combat economic downturns. These programs seek to incentivize lenders to offer loans to riskier borrowers at lower prices. However, the pass-through of guarantees could be heterogeneous across the distribution of risk, and the price changes may be amplified or offset by the effects of bank moral hazard. If information about a borrower is costly to acquire, lenders respond to government guarantees by decreasing their precision of information. This adjustment to the information structure disproportionately harms low-risk borrowers, as information that is less precise offsets the gains that borrowers would have otherwise realized.

In this paper, I first disentangle the effect of the guarantee pass-through and the effect of a change to lenders' information quality across the distribution of borrowers. Conditional on the precision of information (i.e., isolating the guarantee pass-through), a move from a guarantee rate of 90% to a rate of 50% leads to 0.3% higher prices, and 1.9% lower borrower surplus, for high-risk borrowers. In contrast, low-risk borrowers receive 0.1% lower price offers and 0.3% higher borrower surplus. Changes to information precision amplify both of these effects, and, in total, surplus accruing to high-risk borrowers declines by 6.6% and that accruing to low-risk borrowers increases by 2.5%.

The distributional consequences of guarantee programs, and their root in the bank moral hazard problem, suggest that room exists in policy design to provide credit-expansionary incentives while limiting the effects of moral hazard. I show that a hybrid policy including a guarantee and a subsidy yields borrower-surplus gains over the baseline of a guarantee rate of 90% and no subsidy. Specifically, a rule with a guarantee rate of 50% and a subsidy set such that expected spending remains constant confers benefits to low-risk borrowers which, in aggregate, outweigh losses faced by their high-risk counterparts. In total, the hybrid policy yields a borrower-surplus gain of between 1.5 and 2%.

Taken together, these results demonstrate that bank moral hazard has a sizeable impact on the equilibrium response to policy changes in small-business lending markets. In the case of guarantee programs, it amplifies the heterogeneous impact of these interventions across the distribution of borrower risk. Furthermore, the counterfactual results suggest that alternative policies could temper the distributional impact by limiting the effects of moral hazard, and could lead to gains in surplus for the average borrower.

Chapter 2

Incentive Structures and Borrower Composition in the Paycheck Protection Program[1]

2.1 Introduction

Public policy is frequently implemented through private actors (e.g., Affordable Care Act Exchanges, Small Business Administration Loan Guarantee Programs). In such situations, the government faces a tradeoff between inducing the private actors to participate and achieving policy goals. To that end, policymakers typically leverage direct subsidies and other forms of incentives to influence the behavior of participants on both sides of the market. Understanding the responsiveness of actors to these incentives is key in policy design.

In this paper, we study the incentive design of the Paycheck Protection Program (PPP). The PPP was a $953 billion loan forgiveness program that aimed to assist small businesses in

[1]This chapter is joint work with Paul Kim. We would like to thank Vivek Bhattacharya, Gaston Illanes, Robert Porter, and Mar Reguant, as well as participants at the Northwestern IO Student Seminar, for helpful comments and suggestions.

keeping their employees on payroll during the Coronavirus pandemic.[2] Loans were issued by banks, but borrowers were eligible for forgiveness if they used a sufficient share of the loan to support payroll expenses. Should the borrower not meet this threshold, a portion of the loan was forgiven and the remaining balance must be repaid at an interest rate and loan term set by the government.[3] A large amount of recent work has analyzed the targeting of the PPP (see, e.g., Granja et al. (2021) and Bartik et al. (2020)) and its impact on employment and business survival (Hubbard and Strain (2020)). Joaquim and Netto (2021) study lender incentives in the program, but, to the best of our knowledge, our work is the first to analyze the role of the PPP's policy design in determining the allocation of loans and the subsequent decision of whether to use the funds on payroll expenses.

In the PPP, when lenders issued a loan, they received a subsidy equal to a pre-specified percentage of the loan amount. Outside of the subsidy, there were two other categories of incentives that influenced lenders' decisions. First, the SBA set the standard for loan forgiveness by fixing a minimum share of the loan amount that must be used on payroll expenses. Second, it set the interest rate and maturity of non-forgiven loans. These interventions indirectly induced participation of lenders by influencing the behavior of borrowers. The presence of both sets of incentives in the PPP makes it a fruitful environment in which to study the design of policies implemented through the private sector. In particular, we seek to address the following question: are subsidies or borrower-side incentives more effective in (1) increasing credit access and (2) targeting (i.e., ensuring loans are used for forgivable purposes)? Importantly, we show that the answer to this question depends on a number of primitives underlying both lender and borrower decisions.

The design of the program has received considerable attention from legislators and the me-

[2]For more details on the program, see https://www.sba.gov/funding-programs/loans/coronavirus-relief-options/paycheck-protection-program

[3]For first-draw loans, the interest rate is 1% and the maturity is either two or five years depending on the date of issuance. See https://www.sba.gov/funding-programs/loans/coronavirus-relief-options/paycheck-protection-program/first-draw-ppp-loans.

dia. In particular, politicians have questioned the structure of the ex-ante subsidy. Speaker of the House Nancy Pelosi stated: "We have to take a look at how banks are compensated. They get a higher percentage for a small loan, but if you get 5% on a $50,000 loan, that's a lot less than getting 1% on a $5 million loan."[4] Despite the attention paid in the policy arena, there is limited work in the economics literature on the design of the program. We bridge this gap.

To illustrate the role of the underlying borrower and lender primitives in determining the efficacy of policy design, we develop a model of PPP lending. Borrowers are differentiated by their observable loan amount, their unobserved propensity to use funds for payroll purposes (i.e., the relative return of payroll versus other uses), and their unobserved level of cash on hand. With knowledge of these primitives, the borrowers choose the share of funds to use on payroll, given the forgiveness standards and other incentives specified by the program. Lenders are differentiated by their fixed cost of lending to a borrower (e.g., up-front administrative costs, verifying borrower materials) and their marginal cost of an additional dollar. They decide whether to approve a borrower's application for a loan. The relationship between the loan amount and other primitives underlying the borrowers' loan-use decisions, as well as the lenders' fixed and marginal costs of issuing the loan, determine how the policy design affects access to funds and the targeting of the program.

The model has two key implications. First, we show that increasing subsidies and relaxing forgiveness standards induce lenders to offer loans to borrowers seeking smaller amounts. This result implies that policymakers have multiple levers at their disposal to induce an expansion of credit. Second, the correlation between loan amounts and borrowers' tendency to use funds for forgivable purposes determines the relative efficacy of ex-ante and ex-post incentives in improving the program's targeting.

[4]https://www.marketwatch.com/story/pelosi-suggests-banks-making-loans-in-small-business-program-shouldnt-get-paid-more-for-serving-bigger-companies-2020-04-27

Given that one set of incentives need not always be better than the other, we evaluate the observed design of the PPP. We exploit temporal variation in the stringency of forgiveness standards through the passage of the PPP Flexibility Act to provide descriptive evidence to (1) validate the model's key implications and (2) assess the program's efficacy in targeting funds to borrowers who use them on payroll. This piece of legislation made forgiveness standards less stringent, decreasing the minimum share allocated to payroll from 75% to 60%. While our current research design does not allow us to recover causal estimates of the policy response, we find evidence consistent with the predictions of the model. The average loan amount falls by between 6 and 7% in the period following the implementation of the Flexibility Act, which is consistent with lenders being willing to issue smaller loans when forgiveness is easier to obtain. This decline in loan amounts disproportionately benefits businesses in wholesale and retail trade, sole proprietors, and those in urban areas. This result implies that the credit expansion induced by changes to borrower-side incentives could alter the composition of borrowers who benefit from the public program.

To assess the program's targeting, we analyze the relationship between loan size and borrowers' propensity to use funds for payroll. We find that the marginal borrowers (i.e., those who receive loans in the post period but would not have prior to the legislation change) are more likely to use funds on payroll than the inframarginal borrowers. Again, while we do not attribute a causal interpretation to this result, it suggests that the observed relaxation of forgiveness standards improved the targeting of the program.

In aggregate, we find empirical evidence consistent with the claim that the policy change improved both access and targeting. The improvement in access was particularly beneficial for a set of borrowers who typically seek smaller loans and for whom banks may face larger fixed costs of loan issuance. In ongoing work, we estimate the primitives underlying the borrowers' and lenders' decision problems. The goal of such an analysis is to consider counterfactual policy designs – for example, a policy in which banks are subsidized per loan instead of per dollar lent – which may have different implications for access and targeting.

The remainder of the paper proceeds as follows. Section 2.2 reviews the related literature, Section 2.3 describes the Paycheck Protection Program, Section 2.4 presents the model and discusses the main comparative statics results, Section 2.5 presents the data and empirical analyses, and Section 2.6 concludes.

2.2 Related Literature

There are a growing number of papers that study the Paycheck Protection Program. These can be divided into four main strands.

A large number of papers focus on the casual impact of PPP on labor market outcomes. Autor et al. (2020) and Chetty et al. (2020) look at the unemployment rate of PPP-eligible vs. non-eligible firms around the 500 employee cut-off and find that the PPP did boost employment although their magnitudes differ. Barraza et al. (2020), Faulkender et al. (2021) and Granja et al. (2021) use regional variation in lender composition to estimate impact of PPP on labor market outcomes in different geographic regions and find mixed results. Barraza et al. (2020) finds that during the first month, the PPP reduced unemployment by 1.4%, and Faulkender et al. (2021) finds that 10% increase in eligible payroll covered by the PPP resulted in 1-2% decrease in weekly initial unemployment insurance claims. However, Granja et al. (2021) finds no significant evidence that the PPP had a substantial effect on local employment outcomes. Doniger and Kay (2021) uses the 10-day delay between first and second round of the PPP and finds that regions with one percentage point fewer delayed loans within the event window have lower unemployment rates by over 10 basis points.

Related to the above, several papers study and evaluate PPP more comprehensively, especially focusing on targeting. Humphries et al. (2020) finds that, despite smaller firms showing largest improvements upon receiving PPP loans, they were less aware/less likely to apply and less likely to get approved. Bartik et al. (2020) and Hubbard and Strain (2020) both find that PPP approval on net is beneficial for the small businesses, increasing their chance of survival. Bartik et al. (2020), however, notes that banks are less likely to approve higher-distressed firms, and more likely to approve firms with existing connections, even though PPP loans do not seem to be more effective for that group. Similarly, Granja et al. (2021) finds no evidence that PPP funds flowed to areas more adversely affected by COVID-19 and that banks played a major role in targeting - bank participation in the initial phase depends largely on bank characteristics, which explains spatial differences in loan disbursements. Joaquim and Netto (2021) also finds that, during the first phase of PPP, firms that were less affected by COVID-19 received loans earlier but the opposite is true for the second phase. The paper also builds a model of PPP allocation with firms and banks, and finds that the PPP saved 7.5 million jobs.

A few papers specifically study how role of different banks in distributing PPP loans. Li and Strahan (2020) shows that relationship banks (i.e., banks that typically associated as having close relationships with their borrowers) played a big role in disbursing PPP loans. James et al. (2021) similarly finds that community banks made loans faster and lent more relative to their assets compared to larger banks. Lastly, Erel and Liebersohn (2020) studies the role of FinTech in PPP lending and finds that FinTech is more likely to be used in areas with fewer bank branches, lower incomes, larger minority shares, and places more severely impacted by COVID-19. The paper also estimates that FinTech expanded the supply of PPP credit rather than substituting borrowers away from banks. Finally, Lopez and Spiegel (2021) examine the role of the PPP Liquidity Facility, and find that it, along with the PPP itself, increased the growth rate of small-business lending.

Next, a number of papers study the impact of PPP take-up on firms, focusing on the impact on firm valuation and other frictions that some firms may be facing. Balyuk et al. (2020) focuses on small, publicly listed firms and finds that although firms with PPP funds experience positive valuation effects, many firms end up returning the funds, which is also associated with increase in the firms' valuation. Cororaton and Rosen (2021) looks at set of all PPP-eligible public firms and documents that firm value declines after the PPP loan announcement and increases after some of the firms return their PPP loans. The paper suggests that reputational harm and negative signaling limit public firms' participation in the PPP.

Lastly, there is a significant body of work on racial disparities in PPP recipients. Fairlie and Fossen (2021) and Wang and Zhang (2021) use regional variation in minority population to show that places with higher minority share received disproportionately fewer PPP loans. Howell et al. (2021) uses PPP loan-level data to show that Black-owned businesses were more likely to obtain their loan from a FinTech lender versus a traditional bank. Among banks, smaller banks were much less likely to lend to Black-owned firms. Chernenko and Scharfstein (2021) studies a large sample of Florida restaurants and finds significant racial disparities between Black/Hispanic-owned firms vs. White-owned firms. It also finds evidence that this is driven by bank-lending compared to non-bank lending and that Black-owned businesses are more likely to substitute away from bank-administered PPP loans to SBA-administered EIDL loans.

Much of the above literature focuses on finding the overall (or lack there of) impact of PPP. Our paper differs in that we focus on the PPP program design and how different policy parameters could affect both the allocation and targeting of loans. While some of the papers do note that banks play a major role in targeting and allocating loans, most of the papers do not model or study how changing the program parameters could impact the effectiveness of the program. To our knowledge, we are the first paper to model and study how the program design, in particular the lender subsidy and borrower forgiveness standards, could alter both the borrowers and lenders' behavior, shifting the equilibrium allocation of loans.

2.3 Institutional Background

2.3.1 Program Description

The Paycheck Protection Program was originally established by the CARES Act, which allocated $349 billion of funding for the loan forgiveness program between April 3 and April 16, 2020.[5] Later, the Paycheck Protection Program and Health Care Enhancement Act allocated an additional $320 billion[6] with the first applications accepted on April 27, 2020.[7] These two acts served as the primary funding sources for the first-draw loans.

The program aimed to provide forgivable loans to businesses with 500 or fewer employees worldwide, though this restriction was relaxed for businesses in NAICS 72, Accommodation and Food Services (Bartik et al. (2020)). Businesses were also required to meet specified SBA size standards and have tangible net worth less than or equal to $15 million as of March 27, 2020. Importantly, these businesses were not required to satisfy a "credit elsewhere" test, meaning they may have been able to receive financing from other sources.[8]

As mentioned in Section 2.1, the program consisted of two sets of incentives. First, lenders received subsidies, as a share of the loan amount, for participating in the program. For loans less than or equal to $350 thousand, lenders received a subsidy equal to 5% of the loan amount. The subsidy rate declined for larger loans – 3% for loans between $350 thousand and $2 million and 1% for loans greater than $2 million.[9] Second, the policy stipulated

[5] https://www.npr.org/sections/coronavirus-live-updates/2020/04/16/835958069/small-business-emergency-relief-program-hits-349-billion-cap-in-less-than-2-week

[6] For further details, see https://www.ama-assn.org/delivering-care/public-health/summary-paycheck-protection-program-and-health-care-enhancement-act.

[7] https://fortune.com/2020/04/23/ppp-sba-paycheck-protection-program-loans-applying-round-2-what-to-know-small-business-application-congress-funding/

[8] See https://home.treasury.gov/system/files/136/Paycheck-Protection-Program-Frequently-Asked-Questions.pdf. The credit-elsewhere test is an important feature of other SBA lending programs, including the SBA 7(a) Program, the agency's largest loan guarantee scheme. Therefore, the subset of borrowers who received funding through the PPP differed from the subset who participated in other lending programs.

[9] https://www.sba.gov/sites/default/files/2021-02/Procedural%20Notice%205000-20091%20-%202nd%20Updated%20PPP%20Processing%20Fee%20and%201502%20Reporting-508.pdf.

forgiveness standards. Prior to the passage of the PPP Flexibility Act, forgiveness required the borrower use 75% of loan proceeds on payroll and associated expenses, including gross salary and wages, tips, vacation and sick leave, holiday pay, health insurance, and retirement benefits.[10] Following the policy change, the borrower was required to use only 60% on payroll.

Businesses were given a pre-specified amount of time, called the *covered period*, during which they were required to allocate the funds. This stipulation limited businesses from reallocating funds between profits and loan proceeds to keep the loan open indefinitely. Again, as we detail in the next subsection, the PPP Flexibility Act changed the program's covered period.

To obtain loan forgiveness, the borrower must complete an application, either directly with the SBA or through their lender, and compile documentation to prove a sufficient share of funds was allocated to payroll. This documentation includes bank account statements or reports from third-party payroll services to confirm compensation amounts, as well as payroll tax forms. With this information, the SBA conducts a review and determines how much of the loan to forgive.[11]

2.3.2 PPP Flexibility Act

After the initial rollout of the program, Congress amended the SBA's guidance through the issuance of the PPP Flexibility Act. This policy change occurred in response to a number of perceived shortcomings of the original program's structure. For example, due to public-health guidance, businesses such as restaurants and bars did not anticipate being able to reopen in time to spend a sufficient share of funds during an eight-week covered period.[12]

[10] There were a number of contingencies associated with the forgiveness standards. For loans above $50 thousand, the forgiveness rate decreased if the borrower reduced wages by more than 25%. Also, if borrowers were unable to spend the stipulated amount on payroll but demonstrated good-faith effort in rehiring or reduced employees' hours in response to public-health guidance, they were still eligible for forgiveness.

[11] https://www.sba.gov/funding-programs/loans/covid-19-relief-options/paycheck-protection-program/ppp-loan-forgiveness#section-header-4

[12] https://www.jdsupra.com/legalnews/flexibility-act-significantly-improves-40387/

The legislation relaxed the stringency of forgiveness standards.

The PPP Flexibility Act was passed on June 5, 2020, and this piece of legislation altered a number of borrower incentives. It decreased the minimum payroll share from 75% to 60%, increased the repayment period for the non-forgiven portion of the loan from 2 to 5 years, and increased the covered period from 8 to 24 weeks. This policy change provides us with variation in the generosity of forgiveness, primarily through the changes to the minimum payroll share and covered period, with which we analyze equilibrium responses by both borrowers and lenders.

2.4 Model

We model the PPP lending process in two steps, which capture the main decision problems faced by borrowers and lenders. First, a borrower i and lender j are paired, and, after observing the loan amount, borrower characteristics, and the borrower's report of the share it expects to use for payroll purposes, the lender decides whether to issue a loan to the borrower. If the loan is issued, the borrower then receives a shock to its reported share, which determines its final allocation to payroll.

This model captures the responses of both lenders and borrowers to changes in the subsidy rate, forgiveness standards, and loan characteristics (i.e., interest rate and maturity). From the model, we derive two comparative statics results to illustrate the relative efficacy of the two sets of policy levers in (1) expanding access to PPP loans and (2) targeting the loans to businesses that are likely to use the funds on payroll.

2.4.1 Preliminaries

Suppose borrowers are differentiated in three dimensions: (1) b_i, the loan amount, which is a constant multiple of the previous year's profits and is observable and verifiable by the bank, (2) θ_i, a parameter that determines the relative return on funds used for non-forgivable purposes[13] compared to it being used for forgivable (payroll) purposes, and (3) w_i, the borrower's per-period cash on hand or, equivalently, an unobserved profit shifter. Suppose $\theta_i, b_i \sim H_{\theta,b}$, and assume b_i is common knowledge, while θ_i is known only by the borrower. The level of cash on hand is independent of θ_i and b_i, and $w_i \sim H_w$. Lenders j are differentiated by their fixed cost of issuing a loan to borrower i, c_{ij}.

The policymaker sets a subsidy rate, S, which is a share of the loan amount, and three other incentives: $\underline{f}$, the minimum share of funds used for forgivable purposes, r, the net present value of interest paid on the loan should it not be forgiven, and T, the length of the covered period. If the borrower satisfies the forgiveness criteria, then it receives full loan forgiveness. If it does not satisfy the criteria, the loan is partially forgiven, and the borrower must repay the balance at the stipulated interest rate.[14]

Upon receiving the loan, the borrower has a sum of funds equal to the loan amount plus the per-period cash on hand times T, the length of the covered period. At this point, borrowers choose the share of total funds $(T \cdot w_i + b_i)$ to use for payroll purposes, f_i. Upon obtaining the loan, the borrower then receives a shock to its payroll share, ϵ_i. This shock follows a mean-zero, symmetric distribution: $\epsilon_i \sim F_\epsilon$. Note that the government and the econometrician do not observe this final share. Instead, they observe the reported share of the loan amount (b_i) used for payroll. This is observed both at loan origination, $\tilde{f}_i$, and after repayment, $\tilde{f}_i^{post}$:

$$\tilde{f}_i = \min\left\{1, f_i\left(1 + \frac{T \cdot w_i}{b_i}\right)\right\}$$

[13]To be more precise, non-forgivable purposes means not just non-payroll expenses but funds used for other allowed business expenses that fall under broad categories as defined by PPP.

[14]For more information on the terms of forgiveness, see https://www.sba.gov/funding-programs/loans/coronavirus-relief-options/paycheck-protection-program/ppp-loan-forgiveness#section-header-0.

$$\tilde{f}_i^{post} = \max\left\{\min\left\{1, \tilde{f}_i + \epsilon_i\right\}, 0\right\}$$

If $\tilde{f}_i^{post} \geq \underline{f}$ then all of the initial funds are forgiven, otherwise only a portion of the funds are forgiven. Specifically, for $\tilde{f}_i^{post} < \underline{f}$, the share forgiven is $1 - \frac{\tilde{f}_i^{post}}{\underline{f}}$.

2.4.2 Borrower's Problem

Borrower i, seeking a loan of size b_i, is informed of its propensity to use funds for payroll purposes, θ_i, its cash on hand, w_i, the forgiveness standards of the government, $\underline{f}$, the net present value of all interest paid on the loan should it not be forgiven, r, and the length of the covered period, T. The borrower sets its share of funds used for payroll by solving:

$$\max_{f_i \in [0,1]} \gamma\left(f_i + (1-f_i)R(f_i, \theta_i)\right)(b_i + T \cdot w_i) - \\ \mathbb{1}\left\{f_i\left(1 + \frac{T \cdot w_i}{b_i}\right) < \underline{f}\right\}(1+r)b_i\left(1 - \frac{f_i}{\underline{f}}\left(1 + \frac{T \cdot w_i}{b_i}\right)\right), \quad (2.1)$$

where $R(\cdot, \cdot)$ is the return on funds used for non-payroll purposes relative to that on funds used for payroll purposes (i.e., $R(\cdot, \cdot) = 1$ means that borrowers are indifferent between using the funds on forgivable and non-forgivable purposes) and γ is the per-dollar value or return on funds used for payroll.

To ensure the borrower's problem has a unique solution, conditional on T and w_i, and that the optimal share assigned to payroll is monotonic in the borrower's type, θ, we make a number of assumptions on the form of $R(\cdot, \cdot)$. In particular, we assume the function is twice differentiable in each of its arguments and impose three further conditions:

Assumption 1. *$R(f, \theta)$ is twice differentiable in both of its arguments with (i) $\frac{\partial R}{\partial \theta} > 0$, (ii) $\frac{\partial R}{\partial f} > 0$, and (iii) $\frac{\partial R^2}{\partial \theta \partial f} < 0$.*

Assumption 2. *$R(f, \theta)$ is concave in its first argument.*

Assumption 1(i) is without loss of generality and imposes monotonicity of R in θ, Assumption 1(ii) assumes decreasing returns to scale, while Assumption 1(iii) imposes a single-crossing condition. Assumption 2 is more restrictive and is sufficient, along with the other assumptions, to guarantee uniqueness and monotonicity. However, it may not be a necessary condition, and work is in progress to weaken this assumption.

Under the above assumptions, the borrower's optimal share f_i^* is weakly decreasing in its type, θ_i. We summarize this result in the following proposition, the proof of which is in Appendix B.1.

Proposition 1. *Under Assumptions 1–2, for each θ_i, (i) there exists a unique $f^*(\theta_i) \in [0,1]$ that solves the borrower's optimization problem, and (ii) $f^*(\theta_i)$ is weakly decreasing in θ_i.*

Figure 2.1 displays the mapping from borrower type to the optimal payroll share given by the solution to the borrower's maximization problem. The plot illustrates the quantity both as a share of total funds (f_i^*) and as a share of the loan amount reported at origination ($\tilde{f}_i^*$). There are four distinct cutoff types at which the mapping discretely changes. These points of non-differentiability signify changes between corner solutions (i.e., $f^* = 1$, $f^* = \underline{f}$, and $f^* = 0$) and interior solutions.

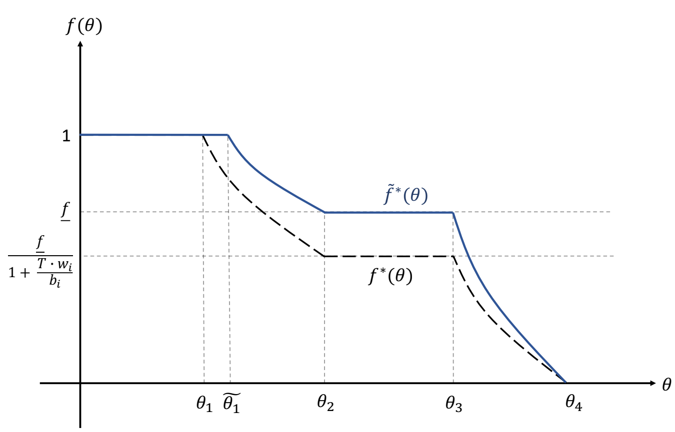

Figure 2.1: Borrower's Optimal Payroll Share

There are two main takeaways from Figure 2.1. First, the result of Proposition 1 is apparent. Conditional on the level of cash on hand, w_i, the optimal payroll share is weakly decreasing in the borrower's type. Second, the level of cash on hand – in particular, its magnitude relative to the loan amount – determines the share of borrower types who choose a payroll share equal to one or a share equal to the policy cutoff. It does not, however, influence the share of borrowers who choose a share of zero. When estimating the model, which is in progress, this type of variation informs the distribution of cash on hand.

The mapping from Figure 2.1 fully characterizes the optimal decision of a given borrower. The lender has knowledge of this rule, as well as a number of borrower observables. We now move to describe the lender's side in more detail and, importantly, its decision of whether to issue a loan to a borrower.

2.4.3 Lender's Problem

Suppose the borrower truthfully reports the payroll share at origination, $\tilde{f}_i^*$, to the lender[15], and the lender observes the loan amount, b_i, and other borrower covariates. The lender does not observe the realization of the borrower's cash on hand, w_i, but knows the distribution from which it is drawn. With this information in hand, lender j decides whether to approve the loan application of borrower i.

The lender's profit from issuing a loan is given by:

$$\pi_{ij} = S(b_i) - c_{ij} + s(\tilde{f}_i^*)(r - \delta_j)b_i,$$

where $S(b_i)$ is the subsidy offered to the lender for issuing a loan of size b_i (in this case, $S(b_i) = Sb_i$), c_{ij}[16] is the borrower-lender specific fixed cost of processing the loan application,

[15] Borrowers are legally constrained from lying on their loan application. The borrower must "certify that the information provided in [the] application and the information provided in all supporting documents and forms is true and accurate in all material respects." False statements are punishable by imprisonment of up to five years or a fine of up to $250,000. See https://www.sba.gov/sites/default/files/2021-03/BorrowerApplication2483ARPrevisions%20%28final%203-18-21%29-508.pdf.

[16] This cost captures, for example, whether the borrower and lender have a pre-existing relationship.

r is the net present value of interest payments, δ_j is the net present value of the marginal cost of lending, and $s(\tilde{f}_i^*)$ is the expected share of the initial loan that is not forgiven and remains a loan. The payoff in the event of not issuing the loan is normalized to zero, so the lender approves the loan if $\pi_{ij} \geq 0$.

Given the structure of the ex-post shock to the payroll share, ϵ_i, the lender's expected share forgiven takes the following form:

$$s(\tilde{f}_i^*) = E_{\epsilon_i}\left[\mathbb{1}\{\tilde{f}_i^{post}(\epsilon_i) < \underline{f}\}\left(1 - \frac{\tilde{f}_i^{post}(\epsilon_i)}{\underline{f}}\right)\right].$$

To ensure the solution to the lender's problem is a unique cutoff rule, we require one further assumption.

Assumption 3. *For all lenders j, (i) $r - \delta_j < 0$ and (ii) $S + s(\tilde{f}_i^*)(r - \delta_j) > 0$.*

The first part of the assumption implies that lenders earn a loss if a loan were not forgiven at the observed interest rate, r. The program stipulates a rate of 1%, and anecdotal evidence suggests lenders prefer loans to be forgiven. If this assumption did not hold, lenders would earn a higher return on non-forgiven loans, given they are fully guaranteed. The second part of the assumption implies the subsidy rate is high enough such that a lender would issue a loan of some amount. Again, this assumption is consistent with the lenders' decisions to participate in the program.

Proposition 2 summarizes the solution to the lender's decision problem and describes a number of comparative statics results. A proof is available in Appendix B.1.

Proposition 2. *Under Assumption 3, there exists a minimum loan amount, $\underline{b}$, such that lender j approves all loans with $b_i > \underline{b}$. The minimum loan amount is: (i) decreasing in S,*

Lenders may face lower costs of loan issuance if the business has a checking account with that bank.

(ii) decreasing in r, and (iii) increasing in $\underline{f}$. Specifically, $\underline{b}$ is given by:

$$\underline{b}(\tilde{f}_i^*) = \frac{c_{ij}}{S + (r - \delta_j)s(\tilde{f}_i^*)}.$$

This proposition illustrates the mechanisms available to the policymaker to induce lenders to extend funds to borrowers seeking smaller amounts. In particular, the policymaker may increase the subsidy rate, increase the interest rate on the non-forgiven portion of the loan, or make forgiveness standards less stringent. However, despite all levers leading to an expansion of lending activity, they have different implications for the targeting of the program. In the next subsection, we illustrate two comparative statics results that show that the relative efficacy of changes to ex-ante subsidies and ex-post interventions (i.e., interest-rate or forgiveness-standard changes) in program targeting depends critically on the relationship between loan amount and the borrower's propensity to use funds on payroll.

2.4.4 Policy Design

To examine the relative efficacy of the policies available to the regulator, we first define its objective. We do not take a stance on the social welfare function – instead, we examine comparative statics on one metric of use for the regulator, the average share forgiven. Because the stated aim of the PPP is to provide businesses with the funding required to keep employees on payroll, we consider this metric a measure of how well the program is targeted. Specifically, define the average payroll share as:

$$PS = \frac{\int_{\underline{b}}^{\infty} b_i \tilde{f}^*(\theta_i)\, dH_{\theta_i, b_i}}{\int_{\underline{b}}^{\infty} b_i \, dH_{b_i}}.$$

With a slight abuse of notation, we do not explicitly consider the relationship between $\tilde{f}^*$ and w_i. Because $\tilde{f}^*$ is monotonic in θ_i for all w_i and the comparative statics results hold for each w_i, then they hold when integrating over w_i.

With the objective function in hand, we first consider the impact of the ex-ante subsidy on the average payroll share. From the previous subsection, we know this policy lever is effective in expanding access to PPP funds. Now, we show that making subsidies more generous has an ambiguous effect on the program's targeting. Depending on the primitives underlying the borrowers' decisions, this change could either increase or decrease the average share of funds allocated to payroll.

Consider a change in the subsidy rate from S' to S'' where $S'' > S'$. The change in the average payroll share takes the form:

$$PS(S'') - PS(S') =$$
$$\frac{\int_{\underline{b}(S'')}^{\underline{b}(S')} b_i \, dH_{b_i}}{\int_{\underline{b}(S'')}^{\infty} b_i \, dH_{b_i}} \left(\underbrace{\frac{\int_{\underline{b}(S'')}^{\underline{b}(S')} b_i \tilde{f}^*(\theta_i) \, dH_{\theta_i, b_i}}{\int_{\underline{b}(S'')}^{\underline{b}(S')} b_i \, dH_{b_i}}}_{\text{Forgiveness rate of marginal borrower}} - \underbrace{\frac{\int_{\underline{b}(S')}^{\infty} b_i \tilde{f}^*(\theta_i) \, dH_{\theta_i, b_i}}{\int_{\underline{b}(S')}^{\infty} b_i \, dH_{b_i}}}_{\text{Forgiveness rate of average borrower}} \right).$$

The two forces determining the aggregate impact of a change in the subsidy rate are apparent in the above expression. Whether an increase in the generosity of the subsidy leads to more funds allocated to payroll depends on the relative strength of the contribution of the marginal borrower and that of the average borrower. If the marginal borrower is more likely than the average borrower to use funds for payroll, then a larger subsidy leads to a greater average share allocated to payroll. The opposite is true if the marginal borrower has a lower propensity to use funds on payroll than the average borrower. The primitive relationship between the borrowers' propensity to use funds for payroll and the loan amount determines which of the two forces dominates.

The effect of a change to the stringency of the forgiveness standards instead depends not only on the primitive relationship between loan amounts and the borrowers' propensity to use funds on payroll but also on how responsive inframarginal borrowers are to changes in the forgiveness rules. Consider a move from a threshold of $\underline{f}'$ to $\underline{f}''$ where $\underline{f}'' < \underline{f}'$. In this case, the change to the average payroll share is:

$$
\begin{aligned}
&PS(\underline{f}'') - PS(\underline{f}') = \\
&\qquad \frac{\int_{\underline{b}(\underline{f}'')}^{\underline{b}(\underline{f}')} b_i \, dH_{b_i}}{\int_{\underline{b}(\underline{f}'')}^{\infty} b_i \, dH_{b_i}} \left(\underbrace{\frac{\int_{\underline{b}(\underline{f}'')}^{\underline{b}(\underline{f}')} b_i \tilde{f}^*(\theta_i, \underline{f}'') \, dH_{\theta_i, b_i}}{\int_{\underline{b}(\underline{f}'')}^{\underline{b}(\underline{f}')} b_i \, dH_{b_i}}}_{\text{Forgiveness rate of marginal borrower}} - \underbrace{\frac{\int_{\underline{b}(\underline{f}')}^{\infty} b_i \tilde{f}^*(\theta_i, \underline{f}') \, dH_{\theta_i, b_i}}{\int_{\underline{b}(\underline{f}')}^{\infty} b_i \, dH_{b_i}}}_{\text{Forgiveness rate of average borrower}} \right) \\
&\qquad\qquad + \underbrace{\frac{\int_{\underline{b}(\underline{f}')}^{\infty} b_i \left(\tilde{f}^*(\theta_i, \underline{f}'') - \tilde{f}^*(\theta_i, \underline{f}') \right) \, dH_{\theta_i, b_i}}{\int_{\underline{b}(\underline{f}'')}^{\infty} b_i \, dH_{b_i}}}_{\text{Change in forgiveness rate for inframarginal borrowers}} .
\end{aligned}
$$

Given $\underline{f}'' < \underline{f}'$, the final term is necessarily negative. Inframarginal borrowers have an incentive to weakly decrease their share allocated to payroll, as lower allocations still receive full forgiveness. Thus, the aggregate effect of a decline in the stringency of forgiveness standards on payroll share is negative if the average borrower is more likely than the marginal borrower to use funds on payroll. If, instead, the average borrower has a lower propensity to allocate funds to payroll, then the aggregate impact of the policy change is ambiguous and depends on the relative strength of the three components described in the expression above. Importantly, in contrast to the case of a subsidy, the adjustment in behavior of inframarginal borrowers can play a pivotal role.

The comparative statics results described in the subsections above provide one key testable implication of the model – the monotonicity of $\underline{b}$ in ex-ante subsidies and forgiveness standards. Furthermore, the model suggests that the correlation of the loan amounts and the share allocated to payroll is of first-order importance when judging the impact of policy changes on program targeting. In Section 2.5, we test the main implication of the model using policy variation from the PPP Flexibility Act. We then provide evidence of the primitive correlation between the loan amounts and payroll shares, showing the observed decrease in forgiveness stringency led to more funds allocated to payroll.

In ongoing work, we plan to take this model to data, estimating the primitives underlying the borrowers' decisions. Using these estimates, we plan to recover the fixed costs of lending. With the primitives, we can simulate counterfactual policy designs to characterize optimal policies for different government objectives, including the objective we have stressed in this section.

2.5 Descriptive Results

2.5.1 Data

For our empirical analysis, we rely on loan-level data from the U.S. Small Business Administration (SBA), which maintains a public dataset containing each approved PPP loan. The dataset contains loan characteristics, including the loan amount, the date the loan was funded, and the loan term. It also contains the borrower's name and address, as well as characteristics such as the NAICS code of the business, a coarse indicator of the business age, and the business type (i.e., individual, corporation, etc.). Finally, it provides the name of the originating and servicing lender.

Outside of characteristics, we observe measures of expected loan use and ex-post loan performance. Specifically, the dataset lists, at origination, the expected share of loan proceeds allocated to utilities, payroll, mortgage interest, and rent, among other uses. The ex-post outcome of interest is the total amount forgiven, from which we calculate the ex-post forgiveness share.

We augment the loan-level data with bank balance sheet information from the Federal Financial Institutions Examination Council's (FFIEC) Uniform Bank Performance Reports.

We use data from March 31, 2020 and match using the bank name. From these reports, we obtain the banks' Tier I Leverage Ratio as a proxy for the shadow cost of lending.

Because our empirical analysis focuses on lenders' responses to the PPP Flexibility Act, we restrict attention to a small window around the policy change. We consider loans issued up to four weeks before the event and up to eight weeks after the event. We analyze lending decisions within this window for two main reasons. First, the composition of borrowers who received loans in the first four weeks after the implementation of the PPP (which occurred eight weeks before the passage of the PPP Flexibility Act) differed from the composition who received them later on. To isolate the impact of the policy change, we seek to standardize the set of borrowers. Second, the take-up of the policy change occurred over time. A number of loans issued in the first four weeks after the policy change indicated an expected payroll share of 75%, the pre-period threshold. The share of loans issued with this expected share declines to close to zero by the fifth week after the policy range. Therefore, our window is not symmetric, and we include loans issued up to eight weeks after the change. We provide empirical support for our window definition in Appendix B.3.

We make a few further restrictions to isolate the equilibrium response to the change in forgiveness standards. We restrict to loans of up to \$350,000. All loans in this category received the same ex-ante subsidy of 5%. This restriction does not eliminate a large number of loans, as only 4.3% of issuances in our twelve-week window are larger than this amount.

Table 2.1 displays summary statistics for our window of interest – twelve weeks around the policy change. There is considerable variation in the size of loans issued. The average loan amount is \$22.7 thousand; however, almost 75% of loans are issued for less than \$20 thousand. This heterogeneity is mirrored in the number of jobs supported, ranging from a minimum of zero to a maximum of 500. The PPP funds support businesses of different ages and geographies – 84% of loans are issued to urban businesses and 83% are issued to businesses that have existed for at least 2 years. Lastly, we observe the share of the loan

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Loan Amount ($000)	1,013,491	22.7	35.1	0.1	5.8	20.8	350.0
# Jobs	1,013,491	3.5	9.0	0	1	3	500
Urban	1,013,491	0.84	-	-	-	-	-
Non-Profit	1,013,491	0.02	-	-	-	-	-
Existing Business (> 2 Years)	833,451	0.83	-	-	-	-	-
Payroll Share (At Orig.)	1,013,491	0.97	0.10	0	1	1	1
Forgiveness Share (Ex-Post)	556,822	0.99	0.07	0	1	1	1

Table 2.1: Descriptive Stats on Approved Loans - 12 Weeks Around Policy

Note: This table presents descriptive statistics for the main analysis sample. This sample includes all loans up to $350,000 issued up to four weeks before and eight weeks after the passage of the PPP Flexibility Act. The sample size differs for two variables due to missing data. In the regressions, when we include fixed effects for business age, we include a category for unknown rather than dropping these loans.

committed to payroll at origination and the share of the loan forgiven ex-post. These final two distributions suggest that forgiveness is common, as the majority of loans are committed to payroll at origination and end up being completely forgiven.

2.5.2 Empirical Tests - PPP Flexibility Act

To recap, our model yields three empirical implications, two that follow directly from the structure of the model, and a third that evaluates whether the observed policy change (i.e., the PPP Flexibility Act) improved program targeting.

Empirical Implication 1. *The minimum loan amount, $\underline{b}$, decreases when forgiveness standards become more generous. Thus, average loan amounts decline in response to the policy change.*

Empirical Implication 2. *The response in $\underline{b}$ is stronger for lenders who face a higher fixed cost (or higher marginal cost) of loan issuance.*

Empirical Implication 3. *Whether the response to the policy change improves program targeting (i.e., the share allocated to payroll) depends on (1) the relationship between bor-*

rowers' propensity to use funds for payroll and the loan amount, and (2) the response of inframarginal borrowers to the policy change.

In the subsections below, we describe the ways in which we empirically test the above predictions of our model. While our research design does not lend itself to recovering causal estimates of the policy impact, the results we present provide support for the structure of the model and suggestive evidence of the efficacy of the policy change. Furthermore, we highlight how the lenders' responses to the policy change alter the composition of borrowers who receive funding under the PPP. In total, these results have important implications for equality in credit access across (1) demographic groups and (2) business types.

Aggregate Response in Loan Size

We begin by considering Testable Implication 1, which implies that lenders respond to the decline in the stringency of forgiveness rules by issuing smaller loans. We first analyze the equilibrium impact of the policy change, testing whether the average loan amount falls in the period following the passage of the PPP Flexibility Act. We estimate event-study specifications of the following form:

$$\log(b_{ijt}) = \alpha \mathbb{I}(t = Post) + \beta X_{ijt} + \epsilon_{ijt}, \tag{2.2}$$

where b_{ijt} is the size of the loan issued to borrower i by lender j in period t and X_{ijt} is a vector of controls, including fixed effects for lender, two-digit NAICS, business type (sole proprietorship, corporation, etc.), borrower state, urban/rural, and business age. Note that the results of this specification could be confounded by underlying time trends in loan amounts. Figure 2.2 displays the path of the average loan amount over time, and this plot suggests the absence of a time trend in the four weeks preceding the policy change. That being said, the pre-period consists of only four weeks, so further support is necessary. The analysis in the next subsection provides further suggestive evidence that a time trend is not the only cause of the observed loan-amount response, and work is in progress to improve the research design to address this shortcoming.

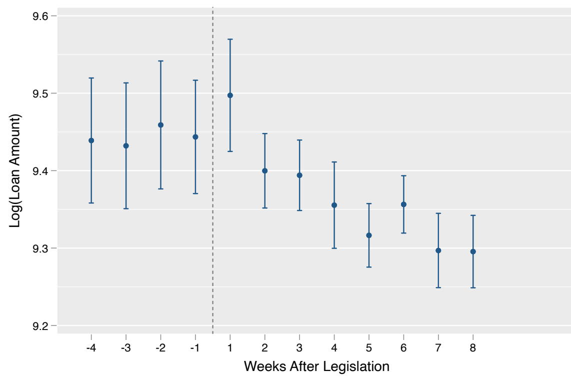

Figure 2.2: Log(Loan Amount) Across Time – 12 Week Window

Note: This figure displays the time trend Log(Loan Amount). In particular, it plots the coefficient of a regression of Log(Loan Amount) on an indicator for the week of loan issuance. The bars denote the 95% confidence interval of the estimate, calculated using standard errors clustered by borrower state.

Table 2.2 displays results for the event-study specifications. The first column displays results with no controls, the second column includes lender fixed effects, while the final column adds the remaining borrower controls. Following the passage of the PPP flexibility act, the average loan amount is approximately 7.2% lower than in the baseline. A change in lender composition explains approximately 15% of this decline. The remainder is explained, almost completely, by a change in the composition of borrowers who receive loans. When we include the full set of controls, the loan amount is only 2.0% lower following the policy change, and this decline is not statistically significant. This result suggests that increasing the generosity of forgiveness may disproportionately benefit certain types of borrowers, namely those who typically seek smaller loans. Sole proprietors and self-employed individuals are two groups whose share of loans is higher following the policy change than in the baseline.

Figure 2.3 further unpacks the impact of the program across borrower covariates and dis-

	(1) Log(Loan Amount)	(2) Log(Loan Amount)	(3) Log(Loan Amount)
Post-Legislation	-0.0723***	-0.0613***	-0.0197
	(0.0205)	(0.0184)	(0.0179)
Observations	1,013,491	1,013,319	1,013,316
Lender FE	No	Yes	Yes
Borrower Controls	No	No	Yes

* p<0.1, ** p<0.05, *** p<0.01

Table 2.2: Aggregate Loan Amount Changes - 12 Weeks Around Policy

Note: This table presents results for the aggregate event-study specifications defined by Equation (2). Standard errors clustered by borrower state are shown in parentheses. Lender FEs are defined as a combination of a lender name, lender city, and lender state. Borrower controls include fixed effects for the two-digit NAICS code, business type (corporation, LLC, sole proprietorship, other), urban/rural, business age ($>$ 2 years, $\leq$ 2 years, unanswered), and borrower state.

plays the change in the share of loans issued to given types of borrowers in the post-legislation period. Panel (a) displays changes in share by one-digit NAICS. Businesses associated with one-digit NAICS codes 4 and 8 are more likely to receive funding following the policy change than they were prior. The former includes wholesale and retail trade, while the latter includes services such as salons and barbershops. Part of the media response to the PPP Flexibility Act, as described in Section 2.3, centered on expanding credit to businesses in industries most affected by public-health measures. This result suggests the policy change may have succeeded in that aim.

Panel (b) examines the pre- and post-legislation shares by business type. In the period following the legislation, sole proprietors received a larger share of loans than corporations and LLCs. This disproportionate impact likely operates through the loan-size channel. Sole proprietors operate smaller businesses than their counterparts. Because loan sizes were a function of prior-year profits, these businesses were constrained, by the program rules, from obtaining large loans. Thus, increasing lenders' willingness to issue these smaller loans confers disproportionate benefits to sole proprietors.

Similar trends emerge in panels (c) and (d). In panel (c), we examine shares by urban/ru-

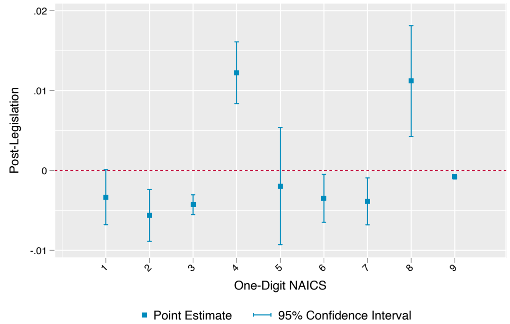

(a) One-Digit NAICS

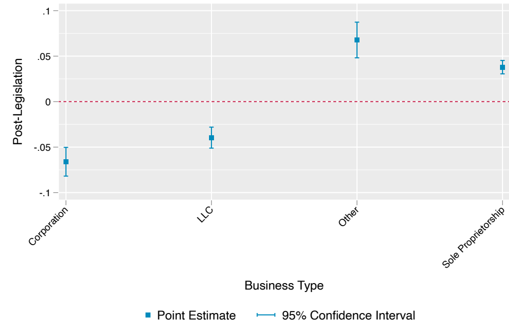

(b) Business Type

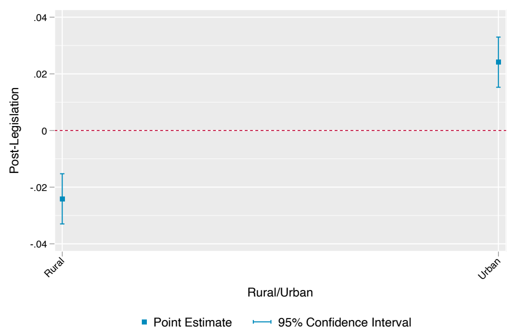

(c) Urban/Rural

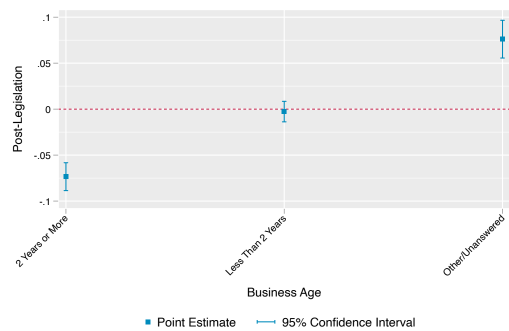

(d) Business Age

Figure 2.3: Heterogeneity in Aggregate Loan Amount Changes

Note: These plots display the coefficient on a regression of an indicator of the specified business characteristic on a dummy for the post-legislation period. The bars indicate the 95% confidence interval calculated using standard errors clustered by borrower state.

ral status, showing that urban borrowers comprise a larger share in the post-legislation period. Finally, panel (d) shows that well-established businesses, those established more than two years prior, are more likely to receive a loan prior to the legislation change. In total, these results suggest that the PPP Forgiveness Act altered the composition of borrowers receiving loans.

It is important to note that these event-study specifications illustrate equilibrium effects of the policy change and, as mentioned previously, could be confounded by underlying time trends in loan amounts. However, taken at face value, this analysis implies a differential impact in the policy change across borrower covariates. In the next subsection, we look instead at heterogeneity across types of lenders to determine whether our model's predictions regarding lending costs hold in the data.

Heterogeneity by Cost

In this subsection, we examine heterogeneity in the policy impact by the cost of loan issuance. This analysis provides evidence in support of Testable Implication 2, which implies that the average loan size decreases by more for lenders who face a high cost of lending. It is important to note that we do not ascribe a causal interpretation to these results. Callaway et al. (2021) show that the standard two-way fixed effects estimator of a generalized difference-in-differences model with continuous treatment does not recover the average causal response under the standard common trends assumption. Instead, the estimate includes a bias term whose magnitude depends on the extent of heterogeneity in treatment effects. Put another way, if the average treatment effect on the treated is more negative for lenders with high issuance costs than for lenders with low costs, then the estimate we recover is downward biased.

In our context, we are most interested in testing whether lenders who find it costlier to issue loans are more responsive to the PPP Flexibility Act. In effect, we seek to pin down the direction, but not necessarily the magnitude, of the average causal response. In this setting, it is reasonable to assume that the average treatment effect on the treated has the same sign for all "doses" of treatment. Under this assumption, the estimate from the standard two-way fixed effects specification has the same sign as the average causal response.

With these caveats in mind, we estimate the following two-way fixed effects specification:

$$\log(b_{ijt}) = \delta_j + \delta_t + \alpha \mathbb{I}(t = Post) \times \tilde{T1}_j + \beta X_{ijt} + \epsilon_{ijt}, \tag{2.3}$$

where b_{ijt} is again the size of the loan issued to borrower i by lender j in period t and X_{ijt} is a vector of controls, including fixed effects for two-digit NAICS, business type (sole proprietorship, corporation, etc.), borrower state, urban/rural, and business age. $\tilde{T1}_j$ is the standardized tier-one leverage ratio for lender j.[17] To remove outliers, we estimate this specification for ratios that fall, inclusively, between the 1st and 99th percentile of the distribution across all lenders.

The tier-one leverage ratio acts as a proxy for a bank's lending cost. This ratio is the sum of tier-one regulatory capital divided by total consolidated assets, and banks must maintain a tier-one leverage ratio of at least 4% to be considered "adequately capitalized" by the banking regulators.[18] PPP loans receive a zero risk weight but are included in the calculation of total consolidated assets unless they are pledged as collateral for a loan from the Federal Reserve's Paycheck Protection Program Liquidity Facility.[19] Thus, banks face a shadow cost of lending due to these regulatory frictions, and that shadow cost is higher for banks closer to the tier-one leverage threshold.

The specification detailed in equation (3) tests whether high-leverage ratio lenders (i.e., those who face a lower shadow cost of lending) are less responsive to the implementation of the PPP Flexibility Act. Table 2.3 displays results for this specification. In the first column, we show results for the specification with no borrower controls, while, in the second column, we add controls, including fixed effects for two-digit NAICS, business type (sole proprietorship, corporation, etc.), borrower state, urban/rural, and business age.

These results illustrate the differential impact of the policy change across bank types. In particular, consistent with the prediction of our model, banks who face a lower shadow cost

[17] To compute the standardized variable, we subtract the mean and divide by the standard deviation across all lenders in the sample. This calculation does not weight by the number of loans issued by a given lender.

[18] https://www.moodysanalytics.com/-/media/article/2011/11-01-03-dodd-frank-act-regulations-minimum-capital-requirements.pdf

[19] For more details, see https://www.elliottdavis.com/ppp-loans-pplf-capital-ratios/. In future work, we plan to exploit the variation across banks in their participation in the PPPLF.

	(1) Log(Loan Amount)	(2) Log(Loan Amount)
Post-Legislation × Standardized T1 Leverage	0.0479***	0.0547***
	(0.0088)	(0.0086)
Observations	743,215	743,214
Borrower Controls	No	Yes

* p<0.1, ** p<0.05, *** p<0.01

Table 2.3: Heterogeneity by Tier 1 Leverage Ratio - 12 Weeks Around Policy

Note: This table presents results for the event-study specifications, allowing for heterogeneity by the lender's tier-one leverage ratio, defined by Equation (3). Standard errors clustered by borrower state are shown in parentheses. Results in both columns include lender FEs are defined as a combination of a lender name, lender city, and lender state. Borrower controls include fixed effects for the two-digit NAICS code, business type (corporation, LLC, sole proprietorship, other), urban/rural, business age ($>$ 2 years, $\leq$ 2 years, unanswered), and borrower state.

of lending (and, therefore, a lower marginal cost of lending) are less responsive to the policy change. A one standard deviation increase in a bank's tier-one leverage ratio is associated with a 4 to 6 percentage point decrease in the change in loan amount between the pre- and post-legislation periods. Given the mean change of between 6 and 7 percent, this is a substantial amount of heterogeneity across lender types. We are in the process of unpacking this result, but, at the very least, it indicates that lender-side heterogeneity should be a consideration when evaluating the efficacy of the policy.

2.5.3 Program Targeting

The preceding empirical results serve two purposes. First, they validate two main empirical implications of the model. Second, they provide evidence of the types of borrowers who benefit from a more generous forgiveness policy. But, we have yet to address whether the policy change improves the program's targeting. In this section, we examine the targeting question through the final empirical implication. In particular, we examine whether the borrowers brought into the program after the policy change are more or less likely to use funds for payroll purposes.

We first estimate regressions to assess changes to the aggregate payroll share before and after the policy change. Specifically, we estimate specifications of the form:

$$f_{ijt} = \alpha \mathbb{I}(t = Post) + \beta X_{ijt} + \epsilon_{ijt}, \tag{2.4}$$

where f_{ijt} is the share of the loan allocated to payroll at origination. The other variables are defined as before.

Table 2.4 displays results for this specification. Following the policy change, the average payroll share is 0.6 percentage points higher than in the pre period. These results suggest that, in aggregate, a decline in the stringency of forgiveness standards is associated with better targeting. It is important to note that the aggregate estimates capture two sets of responses. They combine (1) a change to borrower composition (through lenders' responses in the threshold loan amount, $\underline{b}$) and (2) changes to behavior of inframarginal borrowers. Because we observe a larger share allocated to payroll when forgiveness standards are more lenient, assuming a static distribution of borrowers, the first channel must necessarily dominate. However, we conduct one further analysis to validate the role of this channel, analyzing the relationship between the loan amount and borrowers' tendencies to use funds for payroll purposes.

Figure 2.4 presents a binned scatter plot of payroll share versus the loan amount. In this plot, the negative relationship between these two quantities is apparent. It follows that the marginal borrowers, who receive loans in the post-legislation period, are more likely to use funds on payroll than the inframarginal borrowers. The compositional change leads to an increase in the average propensity, across all borrowers, to use funds for forgivable purposes. Thus, we observe a larger aggregate share dedicated to payroll, as this channel (i.e., the compositional shift) outweighs any adjustment in incentives for the inframarginal borrowers.

There are two important caveats to this analysis. First, as before, the event-study results

	(1) Payroll Share (At Orig.)	(2) Payroll Share (At Orig.)
Post-Legislation	0.0063***	0.0055***
	(0.0011)	(0.0013)
Observations	1,013,491	1,013,483
Borrower Controls	No	Yes

* p<0.1, ** p<0.05, *** p<0.01

Table 2.4: Aggregate Payroll Share Response - 12 Weeks Around Policy

Note: This table presents results for the event-study specifications with payroll share on the lefthand side, defined by Equation (3). Standard errors clustered by borrower state are shown in parentheses. Borrower controls include fixed effects for the two-digit NAICS code, business type (corporation, LLC, sole proprietorship, other), urban/rural, business age ($>$ 2 years, $\leq$ 2 years, unanswered), and borrower state.

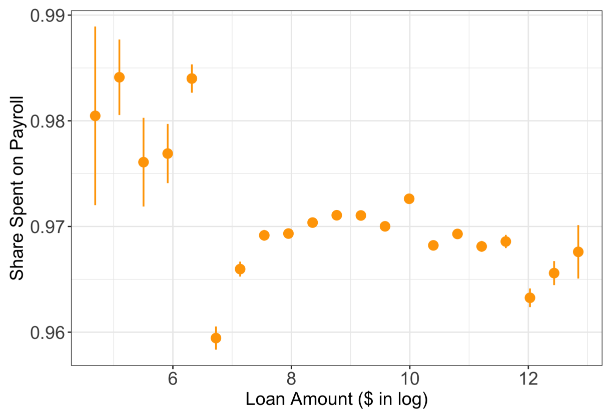

Figure 2.4: Binned Scatter of Payroll Share vs. Log(Loan Amount) – 12 Week Window

Note: This figure displays the binned scatter plot of payroll share vs. loan amount (in log $). Each point represents the average payroll share of observations within each 20 equally spaced bins of log loan amounts. The error bars represent the standard error of the mean of observations within each bin.

could be confounded by underlying time trends in the payroll share. Second, we observe only a local change to borrower incentives. For example, the threshold payroll share decreases from 75% to 60%. Our results suggest that inframarginal borrowers do not respond strongly enough to this local adjustment to outweigh the first channel. However, larger changes to the incentive structure of the program could have different effects. This is one motivation for the estimation of a structural model, which is currently in progress, and the consideration of counterfactual policies.

2.5.4 Discussion

The empirical results, despite their limitations, have a number of implications for the design of the PPP. First, they provide evidence to validate the efficacy of the use of borrower-side incentives (e.g., forgiveness standards and the length of the covered period) in expanding credit to borrowers seeking smaller loans. If lenders prefer that loans be fully forgiven, then easing the burden of achieving forgiveness not only affects borrower decisions but also induces lenders to be more generous. Furthermore, because lending-cost heterogeneity drives a substantial portion of the variation in policy responses, the existence of an institution like the PPP Liquidity Facility can play an important role in the policy design. This setup allows the policymaker to effectively subsidize lenders' marginal costs, whereas the subsidy instead addresses up-front fixed costs of issuance.

The second takeaway from the empirical analysis involves the relationship between loan amount and the borrowers' propensity to use funds on payroll. The observed decline in the stringency of forgiveness standards is associated with an increase in the share of funds allocated to payroll. This change suggests a negative relationship loan size and the propensity to use funds for payroll purposes. The marginal borrowers use funds on payroll, and this change in composition offsets any declines in targeting for the inframarginal borrowers. But, as mentioned above, the relative magnitudes of these two effects may differ under alternative policy designs (e.g., different subsidy rates or larger changes to borrower incentives).

Taken together, these results motivate the use of a structural model to analyze these counterfactual designs. While the evidence in this paper suggests that borrower incentives and marginal lending costs are important drivers of lenders' decisions of whether to issue loans, the descriptive evidence is insufficient to quantify the borrower- and lender-side responses to alternative policies. Estimation of our model is currently in progress, and we aim to quantify the separate impact of fixed-cost subsidies, changes to marginal cost, and forgiveness standards on the composition of borrowers receiving loans through the PPP.

2.6 Conclusion

Quantifying the responsiveness of private actors to the incentives provided by public policies is crucial in determining (1) who participates in the programs, (2) who benefits from them, and (3) how well the programs achieve their targeted aims. In this paper, we analyze the design of the PPP, which aims to provide funds to businesses to keep employees on payroll. We develop a model of PPP lending to illustrate the levers available to policymakers to induce program participation and encourage borrowers to use the loans for their intended purpose. We show that a number of underlying primitives determine which of the levers is most effective in targeting loans to those who seek to use funds for payroll purposes. Importantly, the relationship between the loan amount the borrower is seeking and its propensity to use that fund for forgivable purposes (i.e., payroll) is critical in determining the efficacy of the program.

We then exploit variation in the stringency of the program's forgiveness standards to validate a number of salient features of the model and evaluate whether the policy change improved the targeting of the program. Consistent with our model, we find that average loan amounts are lower in the period following the policy change, which is suggestive of an expansion of lending. This expansion disproportionately accrues to newer businesses, as well as sole proprietors, indicating the importance of policy design in determining the program's impact across demographic groups (e.g., business types). Finally, we find that the program's targeting is better in the period following the decline in forgiveness standards, resulting in an increase in the average propensity to use funds on payroll.

We are in the process of taking the model to data with the goal of considering counterfactual policy designs. While our current framework allows us to analyze a single change to borrower incentives, it does not allow us to consider interactions between fixed-cost subsidies, marginal-cost subsidies, and borrower incentives. The counterfactual exercises will illuminate the tradeoffs faced by policymakers and could guide future policy decisions for programs implemented through private actors.

Chapter 3

Have Mergers Raised Prices? Evidence from U.S. Retail[1]

3.1 Introduction

The Department of Justice and the Federal Trade Commission reviewed over two thousand prospective mergers and acquisitions in fiscal year 2019 (Simons and Delrahim, 2019). The antitrust enforcement agencies are tasked with identifying deals that lessen competition and convincing a court to either block them or force the parties to adopt remedies. This is a difficult task, as the price and consumer welfare effects of mergers are in general ambiguous: the standard academic treatment of horizontal mergers (Williamson, 1968) recognizes that

[1]This chapter is joint work with Vivek Bhattacharya and Gastón Illanes. We are grateful to Igal Hendel, Daniel Hosken, Ariel Pakes, Rob Porter, Mar Reguant, Andrew Sweeting, Bill Rogerson, and Mike Vita for helpful feedback. JD Salas provided excellent research assistance, as did Aisling Chen, Rosario Cisternas, Avner Kreps, Marina Siquiera, and Yintian Zhang. We are also grateful for help from Aaron Banks, Katherine Daehler, Ethan Nourbash, Nathan Friedle, Denis Gribenica, Tianshi Wang, and numerous other research assistants. This project was funded by grants from the Center for Equitable Growth and the National Science Foundation (SES-2116934). Researcher(s) own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. All errors are our own.

increases in market power can be compensated by cost savings due to marginal cost synergies. Furthermore, mergers can induce changes in product quality or lead to repositioning (Sweeting, 2010; Fan, 2013). Accordingly, whether approved mergers typically increase or decrease prices, quantities, and product offerings is an empirical question that is important for evaluating antitrust policy in the US. Researchers have constantly stressed the importance of empirical work on merger retrospectives to understand what mergers have actually done (Whinston, 2007; Carlton, 2009; Ashenfelter et al., 2014).

While a fairly large body of prior work, reviewed in Section 3.2, has conducted such retrospectives, mergers that have been studied are selected on particular dimensions. For a merger to be analyzed in a retrospective, it must satisfy three conditions: (1) the merging parties must have proposed it, (2) the enforcement agencies must have allowed it to go through (or unsuccessfully challenged it), and (3) researchers must have chosen to study it. Each step of this funnel leads to some selection in the set of mergers analyzed. The final step—the decision to even study a merger—is based on dimensions whose effects are unclear: interest in the popular press, data availability, and the potential for publication. Such selection has been shown to lead to significant bias in other economic contexts (Shapiro et al., 2021). Accordingly, even aggregating results over many published studies can lead to an unrepresentative distribution of merger effects.

This paper provides a systematic analysis of the price and quantity effects of mergers in US consumer packaged goods from 2006–2017. We analyze 108 product markets (e.g., canned soup or soluble coffee) in 40 transactions (e.g., a merger between large food conglomerates), consisting of essentially all transactions with a deal size larger than $340 million with consumer packaged goods products likely to be sold through retail outlets. By analyzing the universe of mergers satisfying a particular deal size cutoff, we address the final step of the selection channel: our set of mergers is necessarily representative of (large) mergers that are proposed and approved in this industry. Our second contribution is to shed light on the second selection channel mentioned above—selection on enforcement. By connecting mergers to enforcement actions by the DOJ and FTC, we explicitly study the selection of mergers into approval (without remedies) and in the process quantify agency preferences.

Our baseline estimates of the effects of mergers rely on comparisons within geographies and within products, before and after the merger is completed, controlling for geography-specific time trends and seasonality. We supplement this analysis by controlling for changes in demographics and input costs, to account for demand- and supply-side characteristics that may have price effects. For over 90% of our mergers, we can also use the prices of products in geographic markets where the merging parties have negligible presence as a control.

Our first set of results, presented in Section 3.4, is that merging parties increase their prices by 0.5% on average in the two years following the merger. However, there is substantial heterogeneity in this distribution: the first quartile of price effects corresponds to a price decrease of 2.5%, and the third quartile corresponds to a price increase of 2.8%. Non-merging parties have a small price decrease of 0.2% on average with a slightly narrower distribution of price changes—with an interquartile range of 3.4 pp. Overall, the average effect of mergers on price changes is near zero, averaging to about a decrease of 0.2%. These estimates are fairly robust across specifications.

We repeat the analysis for quantities and measures of product changes. We document a drop in total quantities sold overall: merging parties see a reduction of about 5.5% in quantities sold and non-merging parties see a reduction of about 3.8%. The effects on quantities, however, are far more heterogeneous than the ones on price—with an interquartile range of almost 17 pp for merging parties and 10 pp for non-merging parties. On average, we find no effect of mergers on product introductions or removals—for either merging or non-merging parties.

What contributes to the heterogeneity in these price effects? In the basic Williamson (1968) framework, price changes are due to synergies and changes in market power. Beyond this framework, mergers could induce changes in quality, product assortment, or pricing strategies (such as discrimination). A full welfare analysis of the trade-offs requires a structural model of demand and cost which is beyond the scope of this paper. Instead, we examine patterns in these price changes further to determine the extent to which these channels affect prices.

We first analyze the timing of the price changes. We find that for mergers that led to price increases (in the top quartile of price changes), these price increases materialized very soon after completion: by about 4–6 months after completion, prices stabilized to the new level. This finding is consistent with price increases being due to exercising increased market power rather than quality improvements, which likely take more time. On the other hand, price decreases happen more gradually. Moreover, mergers generally seem to exhibit price decreases between 12–24 months after completion. These latter observations are consistent with the presence of synergies, which take time to materialize (Focarelli and Panetta, 2003; Whinston, 2007).

We continue this analysis by next looking for direct evidence of synergies. We collect data on the production facilities of all firms and use the changes in distance to the nearest production facility as one measure of the synergies generated by the merger. We find that markets with greater decreases in distance typically have larger price decreases. This is consistent with mergers engendering distribution synergies. Although this is merely one source of these synergies, systematic evidence of synergies has remained limited in the literature.

We also study the correlation between price changes and market structure. Agencies often use measures of market concentration—such as the Hirschman-Herfindahl Index (HHI) and the change in this index (ΔHHI)—as measures of the competitive impact of mergers: mergers with especially high HHI or ΔHHI are "presumed likely to enhance market power."[2]

[2] The 2010 revision of the Horizontal Merger Guidelines does advocate from a move away from such "structural presumptions" towards broader analysis of the merger.

Using comparisons both across geographies within a merger and across mergers, we find that prices changes tend to be higher when overall ΔHHI is larger (especially for mergers with low HHI), consistent with standard theories (e.g., Nocke and Whinston (2021)). However, we see no similar impact of the post-merger HHI itself. We do not see any significant patterns in terms of HHI and ΔHHI when comparing price changes within-merger across geographic markets.

These results provide a systematic analysis of the effects of completed mergers in one sector, which is a quantity that has been discussed by many papers (Carlton, 2009) but rarely documented. However, as mentioned earlier, even though there is no selection from completion to being considered in our dataset, there is still selection into completion: mergers that would be especially anti-competitive do not get proposed, are successfully challenged and blocked, or go through with divestitures. This leads to two caveats about our results. First, we do not claim that these distributions are representative of the effects of all possible mergers, or all profitable mergers. Rather, these distributions are representative of the effects of observed mergers. Second, as Carlton (2009) illustrates, due to both this selection concern and the fact that the agency has at best a noisy signal of the future price change at the time the merger is proposed, the distribution of realized price changes does not inform whether an agency is too strict (blocking mergers with negative price effects) or too lax (allowing mergers with positive price effects).

Our second contribution, therefore, is to evaluate the stringency of current antitrust policy. To do so, we collect data on enforcement actions for each of the mergers in our dataset. In Section 3.6, we estimate a deliberately simple model of the agencies' decision to propose a remedy for a merger. In this model, the agency receives a noisy signal of the price change of the merger and proposes a remedy if this signal exceeds a threshold. Using data on enforcement decisions together with the estimates of the realized price changes, we estimate that the US antitrust agencies aim to propose remedies for mergers with an average price increase larger than 3.7–5.6%. When considering price changes of merging parties, the threshold is 8.2–8.8%. These thresholds aim to tackle Carlton (2009)'s critique of merger retrospectives head on—and in doing so, inform the current debate about the laxness of antitrust standards in the US (Kwoka, 2014; Scott Morton, 2019; Shapiro, 2021; Nocke and Whinston, 2021).

3.2 Prior Research on Merger Retrospectives

Merger analysis is one of the primary policy applications of industrial organization, and studying the outcomes of mergers that have actually occurred is a key input into our understanding of them. In the late 2000s, Whinston (2007) noted that looking at price effects of actual mergers is "clearly an area that could use more research" (p. 2425), and Carlton (2009) highlighted the need for more data to guide antitrust reform. Since then, there have been a growing number of merger retrospectives in the literature, culminating in Kwoka (2014), which provides an especially valuable survey and meta-analysis as well as a framework for organizing this literature. Farrell et al. (2009) and Hunter et al. (2008) provide other surveys, and Asker and Nocke (2021) discusses retrospectives in the broader context of the theory of mergers and collusion.

One class of merger retrospectives involve in-depth studies of a small handful of mergers, usually still focusing on prices and quantities. Papers have studied airlines (Peters, 2006; Kwoka and Shumilkina, 2010; Luo, 2014; Das, 2019), assorted consumer products (Ashenfelter and Hosken, 2010; Weinberg and Hosken, 2013), appliances (Ashenfelter et al., 2013), beer (Ashenfelter et al., 2015; Miller and Weinberg, 2017), hospitals (Haas-Wilson and Garmon, 2011; Garmon, 2017) and gasoline (Simpson and Taylor, 2008; Lagos, 2018). Relative to these papers (and many more that we do not have space to cite in this section), we consider a much larger set of mergers, which helps search for systematic patterns across a variety of mergers.

Another class of studies analyzes groups of mergers in an industry in one go. Examples include Kim and Singal (1993) on airlines and Focarelli and Panetta (2003) on Italian banks, and a larger set are summarized in Kwoka (2014). Such studies provide a more comprehensive view on mergers in an industry, as they are less subject to the critique that the set of mergers under consideration are selected. (For instance, Kim and Singal (1993) study the price effects of all US airline mergers in 1985–88.) However, as Kwoka (2014) notes, these studies often still only report the average price change of all mergers instead of the distribution, and they do not tie price changes to particulars of the merger (and in particular enforcement actions). Our analysis considers distributions and their relationship to enforcement actions, and we interpret the distribution to quantify the agencies' objectives.

Kwoka (2014) puts together estimates with the goal of developing a broad understanding of the effects of consummated mergers.[3] From the studies of single mergers, he concludes that while there is substantial variation across mergers in their price effects, on average prices increase by 4–5%. While this provides an important benchmark, these estimates must be interpreted with some caution. Kwoka (2014) is careful to analyze the universe of reputable papers on merger retrospectives, but the mergers selected by the authors of the underlying papers are guided by data availability, the importance of the merger, and publication potential. This may lead to mergers being unrepresentative of any particular sector of the economy: as Vita and Osinski (2018) point out, for instance, almost one-quarter of the 42 data points studied in the analysis of individual mergers come from mergers of academic journals. More importantly, we cannot say that this distribution is representative of all completed mergers, as the process of forming this meta-analysis does not deal with the

[3] The Federal Trade Commission manages a bibliography of merger retrospectives at `https://www.ftc.gov/policy/studies/merger-retrospective-program/bibliography`. However, we are not aware of a meta-analysis of this larger set of studies, and such a meta-analysis would still suffer from some of the unavoidable sample selection issues in Kwoka (2014).

selection into publication, whose effect is difficult to interpret. One main contribution of our paper is that by developing a database ourselves and selecting mergers independently of any of the aforementioned considerations, we can avoid such selection.

We conclude by mentioning two recent papers that are especially related to this study. Atalay et al. (2020) have an interesting analysis of the effect of mergers on product availability—an outcome not typically considered in merger retrospectives—using a representative sample of retail mergers. They find that mergers lead to a reduction in product availability on average, with the acquiror dropping distant products of the target. Majerovitz and Yu (2021) also estimates price and revenue effects of mergers in retail, documenting large asymmetries in the sizes of targets and acquirors. We view these studies as complementary: using different control groups and different decision on how to process data, we arrive at similar conclusions for the average price effect of mergers in retail. Our study has a more explicit focus on other effects of mergers (such as quantities and product assortment), investigating the presence of synergies, analyzing the structural presumptions, and understanding the preferences of the enforcement agencies.

3.3 Data and Sample Selection

In this section, we discuss the data sources (Section 3.3.1) and process to select the sample for analysis (Section 3.3.2). We then describe a number of properties of the sample (Section 3.3.3) and provide observations about mergers that are approved by antitrust agencies.

3.3.1 Data Sources

The first steps of any merger retrospective are to identify the merger of interest and find detailed data on prices and quantities (and possibly more). Given our goal of analyzing a representative set of mergers in an industry, we must limit ourselves to industries where mergers can be easily identified and price data is systematically available. We begin with the set of mergers tracked by SDC Platinum from Thompson Reuters, which provides comprehensive information on mergers, acquisitions, and joint ventures. We then restrict to mergers involving manufacturers of products that are sold in groceries and mass merchandisers, for which detailed price and quantity data are available in the NielsenIQ Retail Scanner Dataset.

NielsenIQ describes this dataset as providing "scanner data from 35,000 to 50,000 grocery, drug, mass merchandise, and other stores, covering more than half the total sales volume of US grocery and drug stores and more than 30 percent of all US mass merchandiser sales volume. Data cover the entire United States, divided into 52 major markets." The data cover 2.6–4.5 million UPCs, depending on the year, and include food, nonfood grocery items, health and beauty aids, and select general merchandise. For each UPC, Nielsen provides sales at the store-week level, along with the average price at which the product was sold. Nielsen also provides some information about the product, including a number of product characteristics and a classification of the product into a "module." As discussed in Section 3.3.2, we use these modules to guide market definitions. We have access to this dataset from 2006 to 2018.

After identifying mergers and obtaining prices and quantities at the product level, we must map each product to an owner, which is important for allocating products to merging and non-merging parties as well as for computing measures of market structure. Unfortunately, Nielsen does not provide ownership of each product. We therefore augment the Nielsen dataset with information from Euromonitor Passport, which tracks ownership over time. When needed, we supplement Euromonitor with further internet searches to manually match UPCs to owners. This is a departure from prior research working with NielsenIQ data, as ownership of products is usually obtained by looking at the first six to nine digits of a UPC, which correspond to a product's "company prefix"—a unique identifier of the company that owns the UPC. This approach is problematic when dealing with mergers and acquisitions, as the transfer of company prefixes during an acquisition can take up to a year, and there is no hard and fast rule determining whether company prefixes are transferred from acquirer to target after a partial divestiture.[4] Instead, working with Euromonitor Passport data allows us to build the entire product portfolio of the main players in each product market for the relevant time horizon, including immediately before and after the merger.

To account for demand- and supply-side characteristics that could influence prices, we supplement Nielsen with data from a number of other sources. First, for each merger, we collect a set of variables—in general, measures of input costs—that could shift production costs. We list inputs for products in each merger (e.g., wheat for cereal) and obtain commodity price indices, typically available from FRED.[5] Second, we collect data on demographics, to help control for changes in factors that could affect demand. We use demographic data county-level data from the American Community Survey and aggregate to the DMA level.

Transportation costs are another source of supply-side heterogeneity for which we seek to control. To proxy for these costs, we collect production and/or distribution facility locations for each product and compute the distance between the nearest facility and the DMA in which a product is sold. Unfortunately, there is no data source that systematically lists these facilities. In practice, there are three main ways to obtain this information: (1) Company 10K reports filed with the SEC often list production and distribution facilities, (2) the Food and Drug Administration maintains an inspection registry of food production facilities that are mapped to the companies that own them,[6] and (3) company websites sometimes list production facilities. We conduct a comprehensive search of these three souces to obtain facility locations, when available. If possible, we assign specific products to production facilities and otherwise assume that products can be produced at all facilities.

Finally, for our analysis of enforcement stringency in Section 3.6, we require data on

[4]See Section 1.6 of the GS1 General Specifications, Release 22.0, for complete details.

[5]There are cases in which multiple mergers involve the same product module. We ensure that the same set of cost shifters is used for mergers that share product modules.

[6]This registry can be accessed at `https://www.accessdata.fda.gov/scripts/inspsearch/`.

enforcement actions pursued by the DOJ and FTC. In particular, we recover whether the agencies required divestitures for a given deal to be approved and which product markets within that deal were the subject of scrutiny. We obtain this information from publicly-available case filings, including Complaints and publicly-recorded Decision and Order documents, available on the websites of the DOJ and FTC.[7]

3.3.2 Merger Selection and Market Definition

To form our sample, we aim to identify all mergers for which the two parties competed in at least one market during the at-issue period (i.e., the period spanning 24 months before the merger's announcement to 24 months past the merger's completion date). To do so, we begin with SDC Platinum dataset and restrict to completed deals that took place on or after 2007, where (1) either the target or acquirer is in the United States, (2) the acquirer is not classified as "Investment and Commodity Firms, Dealers, Exchanges," (3) the deal involves SIC codes that satisfy a broad interpretation of retail products, and (4) the deal size is above $340 million. Table C.1 in Appendix C.1 tabulates the SIC codes of both the deals that pass these initial filters and of the deals that comprise our final dataset.

Most of the deals that survive this initial filtering process either involve firms that do not sell retail products, or involve firms that only sell products that are not tracked in the NielsenIQ Scanner Dataset. To identify the deals that are relevant, we analyze each merger's press release, as well as the merging parties' SEC filings for the year prior to the merger, and identify their retail brands, if they have any. We then search for those brands in the Product files of the NielsenIQ Scanner Dataset.

Whenever both the target and the acquirer own brands that are present in these files, the next step is to determine whether there is product market overlap between them. NielsenIQ

[7]For more information on available documents, see https://www.justice.gov/atr/antitrust-case-filings-alpha (DOJ) and https://www.ftc.gov/enforcement/cases-proceedings (FTC).

categorizes products into product groups, broad categories such as “Prepared Foods - Frozen” or “Condiments, Gravies and Sauces”, and product modules, finer subcategories such as “Soup - Frozen - Refrigerated”, “Entrees - Meat - 1 Food - Frozen”, “Barbeque Sauces” or “Sauce Mix - Taco”. Although these categories divide products in a reasonable fashion, they are not designed to represent product markets, and the degree of granularity varies significantly across product groups. For example, the Nuts product group includes as modules “Nuts - Cans”, “Nuts - Jars”, “Nuts - Bags”, and “Nuts - Unshelled”, while the Snacks product group has product modules covering meat snacks, pork rinds, potato chips, puffed cheese snacks, pretzels, and popcorn, among others. Rather than defining product markets as either product groups or product modules, we define markets as groups of product modules based on our industry knowledge. These categorizations represent our best attempt to define product market overlap. For example, we group all nuts into a single market, but separate the aforementioned snacks into separate markets. We believe that this is a better approach than uniformly following Nielsen product modules or groups. Table C.2 in Appendix C.1 presents a list of product markets for the deals that are considered in our final sample. We find 87 deals, covering 570 product markets, where both target and acquirer sell at least one product in the same product market. In what follows, we will refer to a product market - transaction pair as a merger, so that if companies X and Y merge and they both sell products in product markets 1 and 2, that deal will generate two mergers for our dataset.

The fact that the merging parties sell at least one product in the same product market does not necessarily imply that the deal involves competitors. It could be the case that they sell products in different geographic markets, or at different moments in time, or one of the two parties could own a UPC with a negligible market share. For a merger to qualify for our final dataset, we require that both target and acquirer own at least one brand in a particular product market, and that they sell the brand in the same geographic market at the same time in a two year window around the completion date of the deal. To check this condition, we look at all UPCs that belong in the product market and that are sold within a two year window of the deal, select those that have a non-negligible market share, and assign each of them to their owners using Euromonitor Passport data.

To determine which UPCs have a non-negligible market share, we compute revenue shares at the DMA-month level and begin by considering all UPCs that have a share of at least 1% in any DMA-month in a two year window after the merger. If this is more than 100 UPCs, we keep the 100 best-selling UPCs and any UPCs that have more than a 5% share in any region-month. With this sample of products, we check market coverage: the fraction of sales volume in the product market that is captured by our selected group of UPCs. If the 10th percentile of the distribution of market coverage across DMA-months is smaller than 60%, we repeat this exercise with 200 UPCs. If continues to be the case, we expand the universe to 300 UPCs. Finally, if coverage continues to be too low, we drop the initial share cutoff from 1% to 0.5% and finally to 0.1%. This procedure allows us to work with a tractable number of products but also to expand the set of UPCs that are included in our analysis whenever the product market is particularly varied. Averaging over all mergers in our sample, the average value of the 10th percentile coverage is 73.2%, and the average value of the median coverage is 81.8%. This coverage reassures us that we are capturing the relevant products in each product market.

Having established ownership for each major UPC in the product market, we only keep mergers where both the target and the acquirer sell at least one major UPC in the 24 months prior to completion of the deal. This yields our final dataset of 108 mergers.

3.3.3 Properties of Approved Mergers

Table 3.1 presents summary statistics for our final sample. Each row in this table corresponds to a Nielsen Product Group, which is a coarser categorization that our product market definitions (in Table C.2), but serves to illustrate in which broad product categories the

Product Group Name	N	Product Market Sales (Million USD / yr)	Merging Parties' Revenue Share	HHI	DHHI
All	107	796.4	21.8%	2,689.1	211.9
Candy	3	2901.5	13.2%	2,763.3	111.3
Gum	2	1000.1	43.3%	3,270.9	69.9
Jams, Jellies, Spreads	1	2.7	56.3%	2,455.2	100.6
Juice, Drinks - Canned, Bottled	1	2295.8	16.6%	1,695.2	17.8
Pet Food	3	1461.1	15.7%	3,641.3	67.5
Prepared Food - Ready-To-Serve	3	140.9	8.5%	4,357.7	16.5
Vegetables - Canned	3	59.0	11.9%	3,688.0	9.0
Cereal	2	1258.5	8.6%	2,271.4	26.3
Coffee	2	1201.4	19.3%	2,248.3	8.2
Condiments, Gravies, And Sauces	10	68.5	46.3%	3,116.2	820.7
Pickles, Olives, And Relish	3	57.0	19.5%	2,435.6	27.0
Spices, Seasoning, Extracts	2	64.3	49.8%	3,459.4	141.5
Bread And Baked Goods	14	797.9	23.3%	2,929.1	182.8
Cookies	1	1670.8	0.6%	2,734.5	0.1
Crackers	1	2716.5	14.6%	2,309.7	88.0
Snacks	9	743.7	10.6%	2,494.9	10.2
Baked Goods - Frozen	2	878.3	27.5%	1,701.7	128.6
Pizza / Snacks / Hors D'Oeuvres-Frozen	1	2687.3	35.7%	973.1	570.5
Prepared Foods - Frozen	2	1437.0	2.9%	1,379.8	4.3
Unprepared Meat / Poultry / Seafood-Frozen	1	366.3	10.6%	2,979.8	27.2
Packaged Meats-Deli	7	1434.6	14.8%	1,705.6	30.2
Detergents	2	1253.5	12.2%	2,455.3	154.8
Beer	2	4762.1	30.3%	3,186.9	538.6
Liquor	2	1189.5	8.1%	1,239.7	40.7
Stationery, School Supplies	2	83.4	11.9%	2,642.4	6.3
Cosmetics	10	205.3	23.8%	2,446.9	221.1
Fragrances - Women	1	88.7	7.6%	2,444.2	8.2
Grooming Aids	2	126.8	6.0%	5,705.3	1.5
Hair Care	7	418.4	25.9%	2,447.3	705.5
Men's Toiletries	2	63.2	33.3%	3,070.4	161.7
Skin Care Preparations	4	330.0	18.0%	2,014.6	96.1

Table 3.1: Summary Statistics for the Final Sample of Mergers

mergers are taking place.

Panels (a) and (b) of Figure 3.1 present histograms of average pre-merger HHI and naive ΔHHI. Most mergers have average (across DMA-month) pre-merger HHIs around 2500, with some mergers reaching values over 5000. Most values of ΔHHI are low, close to zero, but several mergers have values over 200. Panel (c) shows a scatter plot of average pre-merger HHI and naive ΔHHI. The mergers with the largest values of ΔHHI tend to have baseline HHI levels around 3000, and mergers in markets with pre-merger HHI above 5000 seem to only be approved when ΔHHI is equal to zero. Panel (d) presents a scatter plot of average yearly sales of the merging parties (in millions of dollars) and naive ΔHHI. Around half of the mergers with ΔHHI over 500 are small, with average yearly sales for the merging parties below $100 million, but several of them feature ΔHHI near 500 and yearly sales around $1 billion. Overall, these patterns are consistent with the selection process that determines whether we observe a merger being consummated. That is, we expect greater antitrust scrutiny on mergers involving large product markets and high values of ΔHHI and pre-merger HHI. Nevertheless, several mergers involving substantial increases in naive ΔHHI have been approved, even in large product markets.

We also find that small mergers with large values of ΔHHI are able to survive antitrust scrutiny when they are embedded in large transactions. Figure 3.2 makes this point through a scatter plot of size of the transaction (in millions of dollars) and naive ΔHHI, where the size of the dots in the plot are proportional to the merging parties' sales in the merger. Recall that a particular deal often features many mergers, as there is overlap in multiple product markets. We find that extremely high values of ΔHHI—for example, above 1000—are only observed in markets with low sales, when they are a part of a considerably larger deal. For example, the point at the top-right corner of the plot is a market with sales around $200 million a year in a deal worth over $30 billion: a deal where no markets with sales above that amount had product market overlap. At the same time, we also observe ΔHHI values around 500 in mergers that seem to be the main market in a deal.

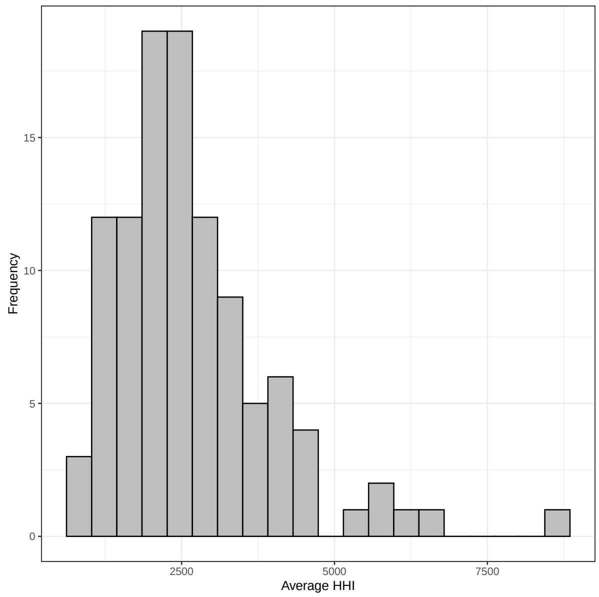

(a) Average Pre-Merger HHI

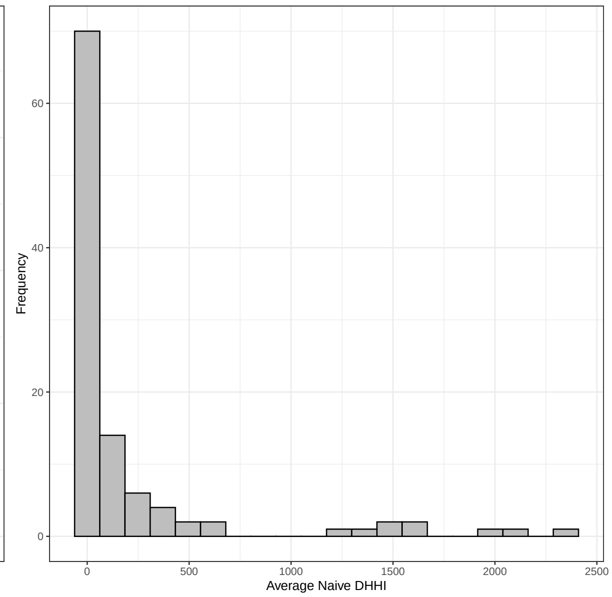

(b) Average Naive ΔHHI

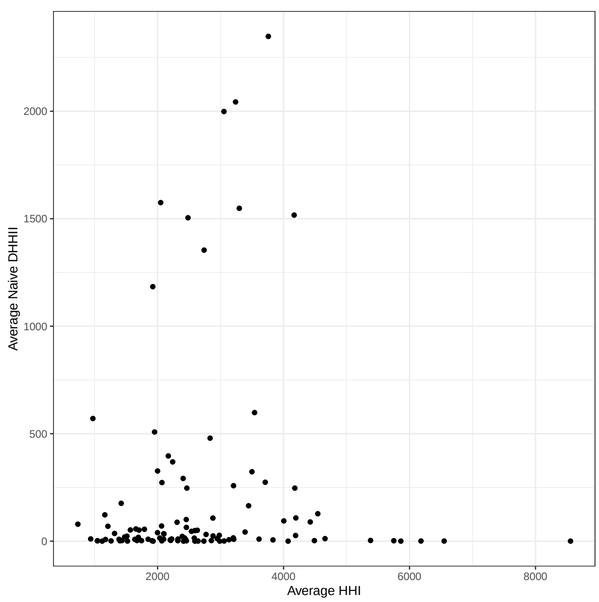

(c) Average Pre-Merger HHI and ΔHHI

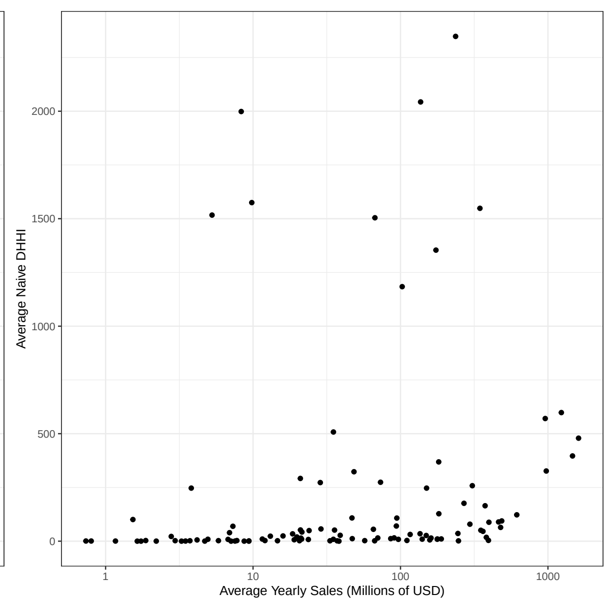

(d) Sales and ΔHHI

Figure 3.1: Distribution of Pre-Merger HHI, Naive ΔHHI, and Merging Parties' Yearly Sales

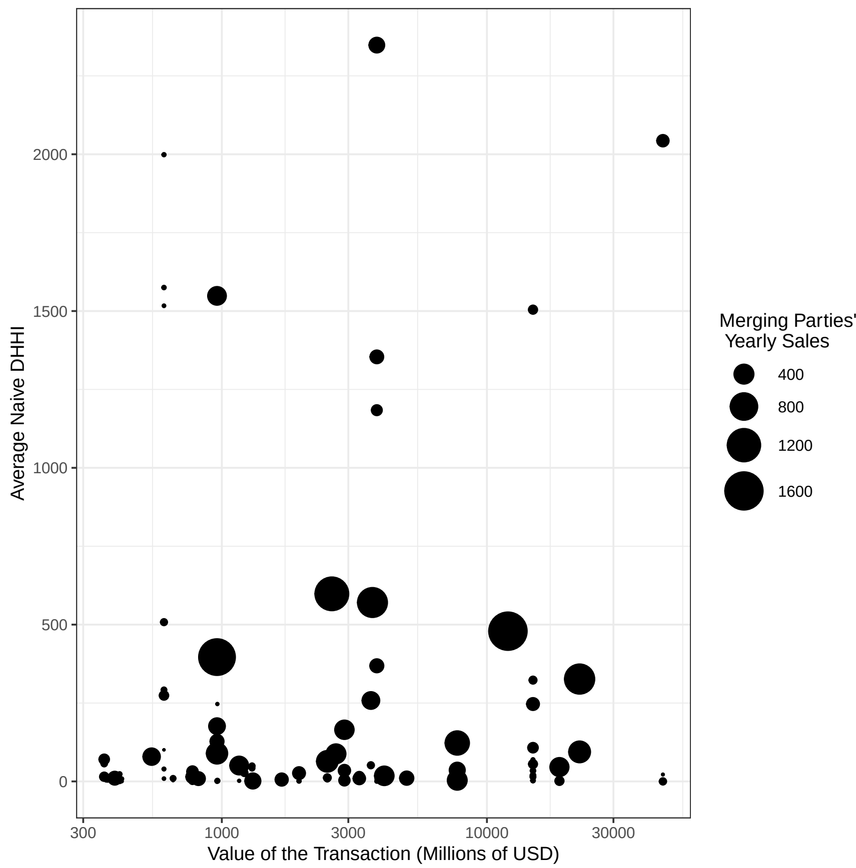

Figure 3.2: Scatter of the Value of the Transaction and Naive ΔHHI, Weighted by Merging Party Sales

3.4 Overall Effects on Prices and Quantities

3.4.1 Empirical Strategy

Retrospective analysis of mergers requires specifying a counterfactual for how prices (or quantities) would have evolved in the absence of the merger. Researchers have typically used one of three classes of controls. The first is the price changes of products of non-merging firms in the same market. For instance, Ashenfelter and Hosken (2010) use private label prices and those of rival products in their study of five consumer packaged goods mergers, and Haas-Wilson and Garmon (2011) use prices of non-merging hospitals in their analysis of hospital mergers on Chicago's North Shore. The rationale is that these products are likely subject to the same cost and demand shocks as products of merging parties. However, an important drawback of this strategy is that non-merging firms are also competitors, and they may adjust their price due to strategic interaction. In this analysis, therefore, we avoid using the prices of non-merging firms entirely.

A second strategy is to use price changes of goods in other markets that are plausibly subject to similar cost and demand shocks. Ashenfelter et al. (2013) study the price effects of the Maytag-Whirlpool merger, which led to consolidation in five appliance markets, by using other appliances not affected by the merger as a control. Kim and Singal (1993) use airline prices in routes that were not impacted by the merger. The advantage of this empirical strategy is that it not subject to the same critique as above: we would likely not expect and price changes in products that are not substitutes for the markets under analysis, and thus any change in prices in such markets are likely due to cost or demand changes. Of course, the difficulty is that since the industry is chosen to be sufficiently distinct, it may not be subject to the same cost and demand shocks as the industries in question. This makes it difficult to find control groups that fit the bill, and it is especially difficult at the scale at which this paper conducts the analysis.

The third strategy is to use price changes in geographic markets where there is no significant change in concentration because one or both parties are small. For instance, Dafny et al. (2012) study the price effects of insurance mergers by effectively using the price changes with low predicted changes in concentration as a control. The rationale for this control group is that in these markets may not have an incentive to adjust price in response to the merger itself. However, these firms may still be subject to national cost and demand shocks that could affect prices. This choice of control group does have two drawbacks. First, not all mergers have geographic markets that are "untreated" by the merger; a priori, this may have been especially problematic in our sample—which contains many large mergers where the merging parties are national staples—although we only lose about 10% of our mergers when defining untreated markets fairly narrowly. Second, this control group does require pricing to be local. If firms price at a national or regional level, as has been documented in some situations in recent work (Adams and Williams, 2019; DellaVigna and Gentzkow, 2019), then a merger could be price "spillovers" onto untreated markets: merging parties may adjust prices in treated markets, causing their competitors to also adjust prices in these markets, which could in turn cause change changes in untreated markets due to national pricing policies.

In this paper, we take two broad approaches to estimating the price effect of mergers. First, we compare prices before and after the merger, controlling for trends and factors that can shift price. To estimate the effect of the merger on outcomes of interest such as average prices and quantities, we regress

$$\begin{aligned}\log p_{idt} = \beta_0 &+ \beta_1 \cdot \mathbb{1}[\text{Merging Party}]_i + \beta_2 \cdot \mathbb{1}[\text{Post-Merger}]_t \cdot \mathbb{1}[\text{Merging Party}]_i \\ &+ \beta_3 \cdot \mathbb{1}[\text{Post-Merger}]_t \cdot \mathbb{1}[\text{Non-Merging Party}]_i + \xi_{id} + \xi_{m(t)} + \text{Controls}_{idt} + \epsilon_{idt}, \quad (3.1)\end{aligned}$$

where i is a UPC, d is a DMA, and t is a month. We discuss the fixed effects and controls below. The coefficients of interest are both β_2, which gives the average difference in prices before and after the merger for the merging parties' products, and β_3, which is the same quantity for non-merging parties. We run variations of (3.1) for other outcomes, and we explain these modifications in the relevant sections below.

In the baseline specification, we add fixed effects ξ_{id} at the UPC-DMA level to allow for persistent price differences in products across geographies. We also include month-of-year fixed effects $\xi_{m(t)}$ to control for seasonality. Finally, we include a DMA-specific linear time trend as a further control to account for secular price changes, such as increases in stocking costs in an area or a persistent decrease in the demand for products. With these controls, we effectively estimate the merger effect as any departure from the trend for pre-merger prices for the same product, in the same geography, at the same time of year: the pre-merger period serves as the control group, and (3.1) can be interpreted as an event study. As a robustness check, we add time-varying controls for demographics at the DMA level (log income per household) and time-varying cost shifters, such as prices of inputs. We do not include any cost shifters that could be directly impacted by the merger: for instance, we do not include proxies for transport costs to the nearest production facility. This check helps us evaluate whether the estimated effects are due to demand or cost structures changing at the time of the merger.

The baseline identification strategy, especially without further cost controls, is broadly based on the idea that any secular trends in demand or cost are gradual; moreover, price data at the monthly level lets us estimate these trends well. Is a linear time trend sufficient to capture any changes in the environment that leads to changes in prices and quantities but not directly tied to the merger? In Section 3.5.1, we show detrended prices for merging and non-merging parties as part of an analysis of the timing of price effects (Figure 3.5): here, we find no further pattern in prices before the merger, after partialling out a DMA-specific linear time trend. Even given this evidence, one may be concerned that the estimated effects are still due to some sort of trend. Two pieces of evidence help us address these concerns. First, we find in Section 3.5.1 that price changes start very soon after a merger: we find it unlikely that the merger completion date would be timed to coincide exactly with when the secular trend in demand or costs would have also led to changes in price. Second, the likely version of this concern is that for mergers in which we document price increases, prices were increasing anyway before the merger—and vice versa for price decreases. However, we find no evidence of pretrends even when splitting the sample by mergers where prices increases significantly, decreased significantly, or did not change much (Figure 3.5 again). Figure C.4 in Appendix C.1 scatters the mean time trend against the estimated price effect, and we do not see any evidence of a positive relationship between these quantities.

Nevertheless, we investigate another identification strategy for robustness. Our second approach uses price changes in geographic markets where the merging parties comprise a small share of total sales—markets that are "untreated"—as a control group. The regression associated with this specification is

$$
\begin{aligned}
\log p_{idt} = \beta_0 + \beta_1 \cdot \mathbb{1}[\text{Merging Party}]_i + \beta_2 \cdot \mathbb{1}[\text{Post-Merger}]_t \\
+ \beta_3 \cdot \mathbb{1}[\text{Post-Merger}]_t \cdot \mathbb{1}[\text{Merging Party}]_i \cdot \mathbb{1}[\text{Treated}]_d \\
+ \beta_4 \cdot \mathbb{1}[\text{Post-Merger}]_t \cdot \mathbb{1}[\text{Non-Merging Party}]_i \cdot \mathbb{1}[\text{Treated}]_d \\
+ \xi_{id} + \xi_{m(t)} + \text{Controls}_{idt} + \epsilon_{idt}, \quad (3.2)
\end{aligned}
$$

where the "Untreated" dummy corresponds to a market with less than 2% market share of both merging parties in the baseline specification. Here, the coefficients of interest are β_3 and β_4. Our definition of untreated markets as having low total presence of merging parties requires these spillovers to happen through non-merging parties. A definition that uses low naive ΔHHI to define untreated markets would also allow for spillovers to happen through merging parties (if only one party is present).

We weigh all regressions by the pre-merger volume at the UPC-DMA level to down-weight the price effect of small UPCs when estimating the average price effect of a merger. When aggregating results across mergers, we weigh estimates by their inverse variance. This practice is recommended by Vita and Osinski (2018) and corresponds to the "fixed effects" aggregation methodology proposed by DerSimonian and Laird (1986) for meta-analysis.

3.4.2 Prices

Table 3.2 presents the distribution of price effects across mergers, both for all products as well as separately for products associated with merging and non-merging parties. The top panel displays results for the baseline specification, the middle adds time-varying cost shifters

	N	Mean	S.D.	25th Pct.	Median	75th Pct.
A. Baseline						
Overall	108	-0.15	3.43	-2.06	-0.74	1.21
		(0.46)		(0.46)	(0.48)	(1.10)
Merging Parties	108	0.48	4.54	-2.54	0.35	2.76
		(0.67)		(0.86)	(0.67)	(1.32)
Non-Merging Parties	108	-0.19	3.37	-2.09	-0.79	1.34
		(0.45)		(0.47)	(0.35)	(0.98)
B. Cost and Demographic Controls						
Overall	106	0.16	2.65	-1.46	-0.18	1.34
		(0.36)		(0.52)	(0.33)	(0.69)
Merging Parties	106	0.82	3.82	-0.65	0.58	2.62
		(0.43)		(1.05)	(0.48)	(1.10)
Non-Merging Parties	106	0.12	2.70	-2.16	-0.19	1.44
		(0.37)		(0.44)	(0.38)	(0.67)
C. Treated/Untreated						
Overall	101	-0.74	4.34	-2.39	-0.35	0.58
		(0.72)		(0.82)	(0.51)	(0.76)
Merging Parties	101	-0.22	5.48	-1.51	-0.18	1.04
		(0.64)		(1.34)	(0.51)	(1.09)
Non-Merging Parties	101	-1.37	3.34	-2.17	-1.62	0.21
		(0.52)		(0.69)	(0.57)	(0.66)

Table 3.2: Overall Price Effects

Note: We aggregate across mergers using the inverse variance of the relevant parameter estimate.

and demographic controls, while the bottom uses markets without merging party presence as a control.

The results from the baseline specification (Panel A) show that mergers on average have no effect on prices: we estimate a decline of about 0.15% for all products. We estimate that merging parties increase prices by 0.48% on average. Non-merging parties have a smaller price change, which we estimate to be a decrease of 0.19% on average.[8] Importantly, in all cases, the standard errors rule out the possibility of large price changes (in either direction) on average. However, these aggregate responses mask substantial heterogeneity across mergers. The standard deviation of price effects is 3.4%. That for non-merging parties in 3.4%, and

[8]Note that the estimates in each row are weighted by the inverse variance of the specific coefficient. Thus, each row in Table 3.2 may have a weighting scheme.

merging parties have a slightly wider distribution of effects, with a standard deviation of 4.5%. The estimated interquartile ranges are comparable to the standard deviation. We find that about 25% of mergers have a price effect for merging parties that exceeds 2.8%.

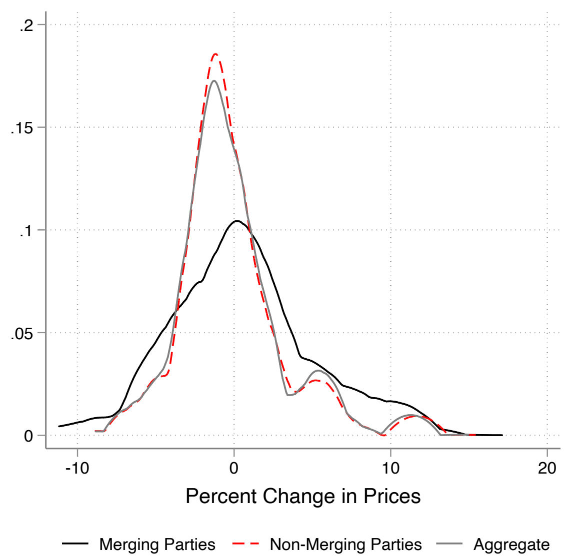

(a) Distribution of price changes

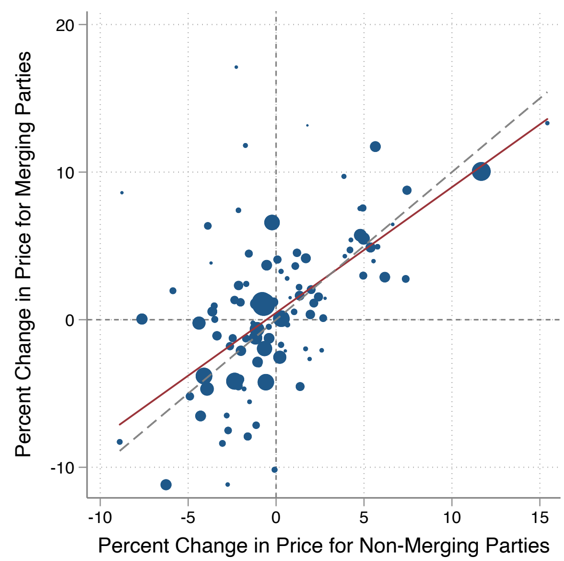

(b) Scatter of price changes

Figure 3.3: Price Changes for Merging and Non-Merging Parties, as Estimated by (3.1)

Note: Distributions in Panel (a) are weighed by the inverse variance of each parameter estimate. The size of each point in Panel (b) corresponds to the inverse variance of the parameter estimate for the merging parties.

These findings are evident in Figure 3.3(a), which plots the distributions of merger effects using the baseline specification. As shown in the summary statistics, we document substantial heterogeneity in price effects both across and within firm types (i.e., merging and non-merging). Panel (b) scatters the estimated effects for merging and non-merging parties against each other; the size of each point corresponds to the inverse variance of the estimate of the merging parties' price change for a given merger. The solid, red line is the best fit of a regression of merging parties' price changes on non-merging parties' price changes, with weights again given by the inverse variance of the estimate of the merging parties' price coefficient. The dashed, gray line is the 45-degree line. As would be predicted by standard models of strategic interactions (i.e., those in which prices are strategic complements), the price changes of non-merging parties are typically of the same sign as those of merging parties—and of a similar but slightly smaller magnitude, as the best fit line is slightly flatter than the 45° line. To our knowledge, there are no general conditions comparing the price changes of merging to those of non-merging parties. A relevant benchmark is provided by Deneckere and Davidson (1985), who show that under fairly general demand systems but ex-ante symmetric firms, merging parties increase prices more than non-merging parties do. While we find this is true on average, it does not hold merger-by-merger in our sample.

Are these estimated price changes due to changes in costs or demand not induced by the merger? Panel B repeats the analysis with controlling for cost and demographic effects. Overall, we find that the point estimates for average price effects are slightly larger (by about 0.4 pp), and the estimated width of the distribution is slightly smaller; an inspection of the quantiles suggests that this is due mostly to a shorter left tail of the distribution of price changes. Nevertheless, the main observations still hold with these estimates: average price changes are small, and there is substantial heterogeneity in the estimated price effects. Moreover, Figure C.1 in Appendix C.1 shows that not only are the estimated distributions of price effects similar but the estimated effects are correlated merger-by-merger.

Panel C uses “untreated” markets as a control, dropping seven mergers in which we have no such markets. At a broad level, the takeaways are similar to the other panels—with a small average treatment effect and a wide distribution overall. The estimated average effects are more negative, with the average overall effects and that for merging parties about 0.6–0.7 pp smaller than in the baseline. The estimates for the nonmerging party is a drop of 1.4%. Interestingly, we find that this is mostly due to a reduction in upper tail of the distribution: the first quartile of the distribution is similar to the one in the baseline, but the third quartile is estimated to the systematically smaller.

One could devise explanations for these results. If price increases are due to industry-wide cost increases or demand increases (reduction in elasticity) not captured in Panel B, then we may expect to see increases in prices in the untreated markets that would call into question whether the price increases are really due to the merger. In Section 3.5.1, we show that price increases materialize quickly after a merger is completed, which seems at odds with (likely more gradual) price and cost changes. Alternatively, the control group itself could be partially contaminated by the merger effect: national or regional pricing policies could lead to price increases in untreated DMAs as well. However, we choose not to read too much into the differences in point estimates between Panels A and C and instead simply view it as a robustness check on the main observations. This is for two reasons. First, the standard error on the point estimates is large, and for both the overall effects and the merging parties' effect, the average from Panel A lies within one standard error of the mean from Panel C. Second, the reduction in the point estimates is sensitive to the choice of weighting scheme to aggregate results: Table C.3 in Appendix C.1 shows results uses the random effects weighting scheme of DerSimonian and Laird (1986), a standard in meta-analysis in biostatistics, and using untreated markets as a control leads to larger point estimates. However, the main observations—that price effects are small on average, they are smaller for non-merging parties, and the distribution of price effects is wide—are robust to all specifications and all weighting schemes.

Taken together, the results in this section contrast with the findings of Kwoka (2014), which documents an average price change of 7.2% (Table 7.2). The difference between our findings and Kwoka's highlights the importance of collecting a comprehensive set of consummated mergers. Our analysis suggests that, on average, the typical approved merger of consumer goods firms leads to modest price increases. However, we document substantial heterogeneity in merger effects. Furthermore, as noted by Carlton (2009) and examined in more detail in Section 3.6 of this paper, average price changes are not indicative of enforcement stringency, which should instead be judged by the expected price effect of the marginal merger.

3.4.3 Quantities

While most merger retrospectives have focused on prices, a natural question is whether mergers have reduced quantities on average. Conventional intuition would suggest that even if mergers have small price effects on average, a significant drop in quantity may be indicative of mergers having adverse welfare effects.[9] While we do not claim that the results reported this section should be interpreted as welfare effects, this intuition would indicate that documenting quantity effects is of direct interest.

To compute quantity effects, we aggregate to either the DMA-month-firm type level, where a firm type is whether that firm is a merging party or not. That is, we sum all quantities of all products sold by all non-merging parties in each DMA-month, and we sum all quantities of all products sold by merging parties in a DMA-month. We regress the log of this quantity on the same right-hand side as (3.1), with DMA-level time trends and DMA fixed effects. We also run a specification where we aggregate to the DMA-month level and simply include a post-merger dummy (again with DMA-specific time trends). This is for two reasons. First, running the regressions at the UPC-DMA level as with price would return an average of how product-level quantities have changed in percentage terms. Given the wide dispersion in quantities for different UPCs (unlike for prices), this average is difficult to interpret. Second, results relating to consumer welfare effects rely on tests of total quantity.

Table 3.3 and Figure 3.4 show results. On average, we find sizable decreases in quantities: about 4% overall, 5.5% for merging parties (although this is driven by a long left tail), and 3.8% for non-merging parties. There is substantial heterogeneity in this estimate, but the magnitude of the drops in quantity at the first quartile is typically substantially larger than

[9]Ongoing work by Lazarev, Nevo, and Town formalizes this intuition: showing that under certain conditions—including an assumption that the welfare effect of a merger has the same sign for all customers—the sign of the effect on total quantity coincides with that of the welfare effect of the merger.

	N	Mean	S.D.	25th Pct.	Median	75th Pct.
Overall	108	-4.00	6.78	-7.91	-4.60	-0.02
		(0.82)		(1.15)	(1.31)	(1.83)
Merging Parties	108	-5.47	14.82	-14.85	0.06	2.06
		(1.93)		(4.12)	(3.86)	(1.91)
Non-Merging Parties	108	-3.75	7.81	-7.07	-3.97	3.06
		(1.19)		(1.67)	(2.01)	(1.59)

Table 3.3: Quantity Effects

the magnitude of the increases in quantity at the third quartile. In fact, we estimate that the third quartile of overall quantity effects is close to 0. Figure 3.4(b) scatters the quantity effects for merging parties against those for non-merging parties, and we generally see a positive association between these estimates.

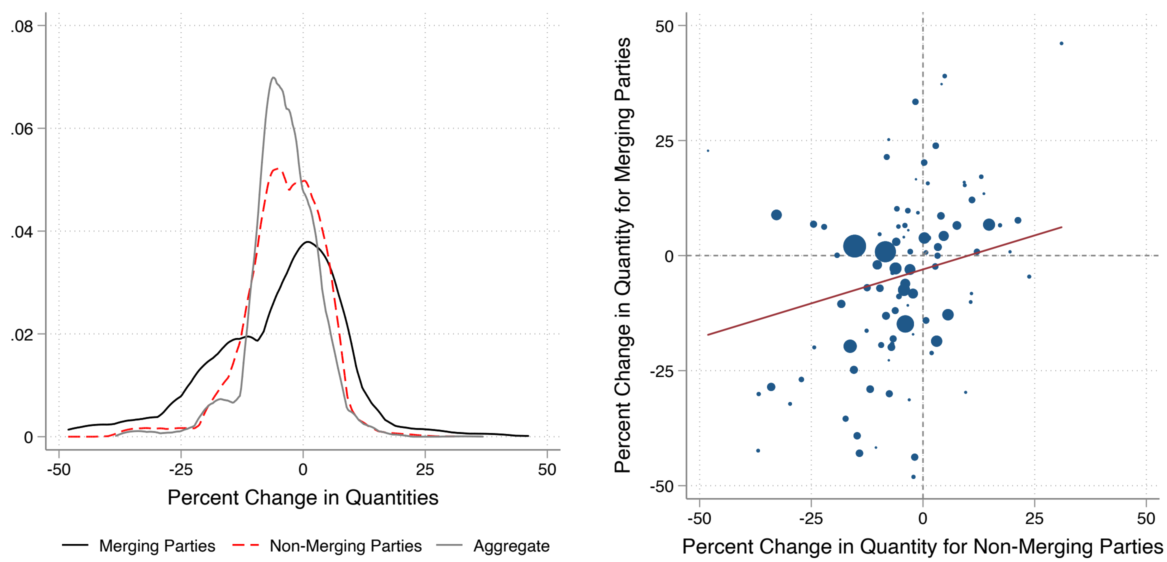

(a) Distribution of quantity changes

(b) Scatter of quantity changes

Figure 3.4: Quantity Changes for Merging and Non-Merging Parties, as Estimated by (3.1)

Note: We restrict to changes that fall in the interval [-50,50]. Distributions in Panel (a) are weighed by the inverse variance of each parameter estimate. The size of each point in Panel (b) corresponds to the inverse variance of the parameter estimate for the merging parties.

3.4.4 Product Assortment and Availability

While merger retrospectives have typically focused on price effects, horizontal mergers can have important effects on product assortment—whether new products are introduced or existing ones are removed. The impact of a merger on these incentives is ambiguous as well. The desire to avoid cannibalization may push merging parties to remove similar products. The cost structure of the new firm also changes, which could lead to effects in either direction: it may make certain product lines unprofitable or make it economical to increase the product line. Finally, the incentive to "innovate" by designing new products changes as well.

Berry and Waldfogel (2001) discuss these competing effects and conclude that the effect of mergers on variety is necessarily an empirical question. While researchers have argued that ignoring product changes after a merger can bias welfare effects (Draganska et al., 2009; Fan, 2013), retrospectives remain rare, with some exceptions in radio (Berry and Waldfogel, 2001; Sweeting, 2010). To our knowledge, Atalay et al. (2020) provide the most comprehensive study on this question, studying product availability in a large set of mergers between retail products. They document a reduction in the number of UPCs offered by merging parties on average. Using a careful analysis of product descriptions and characteristics, they trace this back to firms dropping "fringe" products and focusing on their core competencies. Our first contribution is mostly a replication of their results: we conduct their analysis on a broader sample, albeit with specifications that focus on changes in the product mix rather than the count of products. Second, we also study the geographic distribution of products.

For each merger, we create an original list of products by taking all UPCs sold nationwide in the first six months of our sample (i.e., 18 to 24 months before the merger completion date) which survive our sample selection criteria. For each subsequent month, we compute the proportion of UPCs sold by a firm that are "new"—i.e., that are not included on this original list—computed as the ratio of the number of new products to the number of old products. These results are under "Introduction of Existing Products." We also compute

	N	Mean	S.D.	25th Pct.	Median	75th Pct.
A. National Level						
Introduction of New Products						
Merging Parties	96	0.58	2.91	-0.63	-0.05	0.16
		(0.43)		(0.33)	(0.15)	(0.93)
Non-Merging Parties	106	0.25	0.75	-0.04	0.11	0.26
		(0.06)		(0.03)	(0.06)	(0.19)
Continuation of Existing Products						
Merging Parties	98	0.44	1.32	-0.04	0.11	0.47
		(0.19)		(0.03)	(0.06)	(0.51)
Non-Merging Parties	108	-0.04	0.70	-0.13	0.00	0.03
		(0.05)		(0.11)	(0.04)	(0.04)
B. DMA Level						
Introduction of New Products						
Merging Parties	96	0.50	4.12	-0.66	0.25	1.90
		(0.63)		(0.66)	(0.37)	(0.90)
Non-Merging Parties	106	0.33	1.31	-0.18	0.07	0.48
		(0.22)		(0.29)	(0.06)	(0.38)
Continuation of Existing Products						
Merging Parties	98	3.48	17.13	1.23	1.96	7.25
		(2.20)		(1.47)	(1.49)	(2.37)
Non-Merging Parties	108	-0.81	3.17	-0.86	0.04	0.80
		(0.58)		(1.13)	(0.33)	(0.19)

Table 3.4: Assortment Effects

Note: Effects are in percentage points.

the share of UPCs from this original list that have survived, computed as the number of old products still sold in the month divided by the number of products on the original list, and we list results under "Continuation of Existing Products." To understand regional product introductions and removals, we compute an analogous list at the DMA level as well. We mark a product as new in subsequent months if it is new to the DMA (and mark it as removed if it is removed from the DMA). We use a specification similar to the one in Section 3.4.3.

Table 3.4 shows results. For context, the mean value of introductions of new products across all mergers is 3.45 pp: about 3.5% of the UPCs carried in a particular DMA-month were not in the list at the start of the sample. The distribution of treatment effects for both merging parties and non-merging parties is roughly centered at zero. Thus, like with all other quantities we have studied in this section, we do not see any systematic effect of merging on product introductions or product removals. This is true both for the national-level regressions and the DMA-level ones: not only do we not see evidence that merger systematically lead to legitimately new products, but we do not find systematic evidence that they lead to existing products being sold in a wider range of geographies.[10]

An important observation, however, is that the interquartile range is somewhat large compared to the mean of the variable for merging parties. At the national level, the difference between the third and first quantiles is about one-third the sample mean. Future drafts of this paper will investigate the sources of this heterogeneity. However, a somewhat robust result is that the distribution of effects on non-merging parties is rather tight for any specification: the standard deviation of effects for non-merging parties is at most half that for merging parties—and usually much less. This distribution is also centered around zero. Overall, we conclude that mergers have little effect on the product assortment decisions of non-merging parties.

3.5 A Closer Inspection of Price Changes

The results in Section 3.4 show a wide range of price changes across mergers. What contributes to the heterogeneity in the price effects of mergers? Broadly, the Williamson (1968) trade-off is between increased market power generated by the merger (due to unilateral or coordinated effects) that could lead to price increases, and marginal cost synergies that are passed through to the customer, which lead to price decreases. Beyond the Williamson (1968) trade-off, mergers could also lead to changes in product assortments, changes in quality for the existing set of products, or changes in price discrimination strategies. In this section, we

[10]This observation is a contrast to the ones in Atalay et al. (2020), who find that merger on average seem to reduce the number of products carried by merging parties. We are investigating whether sample selection decisions affect the difference between these two studies.

aim to understand the extent to which these channels impact observed price changes.

A full unpacking of these channels requires estimates of demand and cost functions before and after the merger, which is outside the scope of this analysis. Instead, we rely on providing suggestive evidence that price changes are correlated with measures of these channels. First, we study the timing of price changes, which prior research (Kim and Singal, 1993; Focarelli and Panetta, 2003; Whinston, 2007) has argued could be informative of whether price changes are due to firms exercising market power or realizing synergies. Second, we look for evidence of anticipatory price effects: price changes ocurring after the announcement but prior to completion of the merger. Third, to look for direct measures of synergies, we study whether measures of cost reductions—in our case, changes in distance from the nearest production facility—are related to price reductions. Finally, as a way of evaluating whether firms are exercising the market power generated by the merger, we study the relationship between price changes and measures of concentration.

Why should we be concerned about understanding what drives these price changes? First, we believe that such an analysis is of direct interest. While the Horizontal Merger Guidelines have moved away from structural presumptions, they are still an important summary measure for the agencies: understanding whether realized price changes for completed mergers is correlated with HHI and ΔHHI can inform these presumptions. Moreover, credible possibilities of synergies are a common defense of a merger; however, the literature lacks systematic evidence on whether synergies have actually materialized in many completed mergers. Thus, these analyses can also inform enforcement agencies' decisions by identifying merger characteristics that suggest an increased ability to exercise market power or the potential for cost synergies. Second, while the majority of merger retrospectives have focused on price effects—partly since they are typically the first screen used by agencies—whether mergers have harmed consumer surplus is also a fundamental question of interest. A full analysis of consumer surplus would require a structural demand model that is outside the scope of this paper. However, ruling out the degradation of product quality or decreased product availability would suggest that small price changes are not coupled with significant changes in consumer welfare.

3.5.1 Timing

How quickly after a merger's completion are price changes realized? Conventional wisdom suggests that any strategic effects of the merger due to increased market power would materialize quickly, as it becomes legal to discuss pricing strategy between companies at the date of completion. Thus, if price increases are due primarily to market power, we would expect such increases to materialize soon after the merger. On the other hand, price changes due to adjustments in product quality or synergies would take longer to realize, as the firms integrate their operations and adjust their production and distribution facilities.[11] To evaluate whether changes in market power, product quality, or synergies contribute to the observed price effects, we document systematic evidence on the timing of price changes.

First, we estimate a version of (3.1) with finer bins of time than simply post-merger: we replace the post-merger $\times$ merging and post-merger $\times$ non-merging dummies with months-since-merger fixed effects, interacted with merging and non-merging dummies. We depart from (3.1) in one simple way. We still wish to control for a time trend to control for secular changes in costs and demand, and these fixed effects would be collinear with the time trend. Thus, we first detrend the data with a DMA-specific linear time trend estimated off the pre-period. We also normalize the fixed effects so that the dummies are equal to 0 for the month immediately preceding the merger. The specifications also control for UPC-DMA

[11] Some limited empirical evidence supports this view: Focarelli and Panetta (2003) document that short-run price increases in banking mergers dissipate over the period of 3–6 years, attributing these to synergies materializing. However, researchers have also identified cases where synergies are realized especially quickly: Eliason et al. (2020) show that acquired dialysis firms adapt the parent chain's policies (many of which are harmful for patients) within one or two quarters of acquisition. An exception may be synergies that are due to increased buyer power after the merger, which may materialize more quickly.

fixed effects (but drop seasonality fixed effects, as those are collinear with the months-since-merger effect for each merger). Figure 3.5 displays the mean price change for merging parties, displayed as orange squares, and non-merging parties, displayed as blue circles, across time. We display these plots separately for price changes in the top quartile of the distribution across mergers (high price changes), those in the bottom quartile (low price changes), and those in the second and third quartiles (stable prices).

Our first observation is that, after controlling for a linear time trend, prices tend to be generally flat before the merger is completed. We can interpret this as a response to the natural concern that price effects that we attribute to the merger are instead due to secular time trends. If mergers that we estimate to have large price increases exhibited a trend of price increases before the completion date as well, we would be skeptical of whether the estimated positive price effect is due to the merger itself. Panel (a) shows evidence that this is not the case: price increases happen after the merger, and there is no visible trend on average before the completion date. Panel (b) shows a similar result for mergers that exhibit low price changes (i.e., large price decreases), although there is arguably a small price decrease over approximately 8 months before the completion date itself on average. Panel (c) shows that mergers in the middle 50% of the price change distribution also tend to have a relatively flat price trend.

Of course, the main goal of Figure 3.5 is to illustrate the pattern of price changes after the merger. Focusing on Panel (a), we find that mergers with a substantially large price increase tend to realize that increase fairly quickly. On average, the price increase for merging parties stabilizes within 4–6 months from the date the merger is completed, suggesting that market power effects are realized fairly quickly. We find that non-merging parties also realize a price increase equal to that of merging parties for the first 4–6 months after the merger, but their price increases drop to about half of the maximum level by about 9 months post-merger. We see some evidence that prices of both merging and non-merging parties drop further

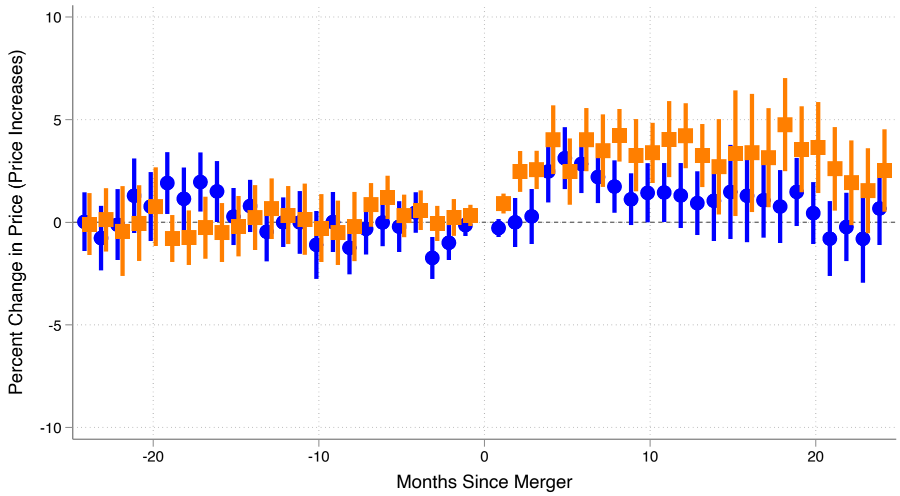

(a) High price changes

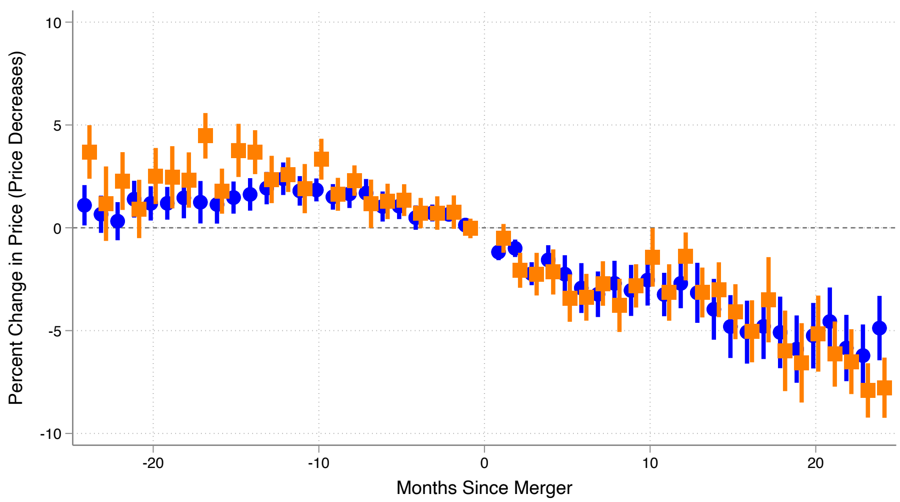

(b) Low price changes

Figure 3.5: Timing of Price Changes, for Merging Parties (orange square) and Non-Merging Parties (blue circle). (Continued on next page.)

Note: The marker indicates the mean price change, relative to a pre-merger DMA-specific, linear price trend, the given number of months after the merger becomes effective, and the thick line is the 90% confidence interval of that mean. Panel (a)–(c) shows subsamples: Panel (a) restricts to mergers with price changes in the top quartile, Panel (b) restricts to mergers with changes in the bottom quartile, while Panel (c) displays the remaining mergers. Panel (d) shows all mergers.

approximately 18 months after the merger, by about 1–2 pp.

The time pattern for mergers with large price decreases, documented in Panel (b), is

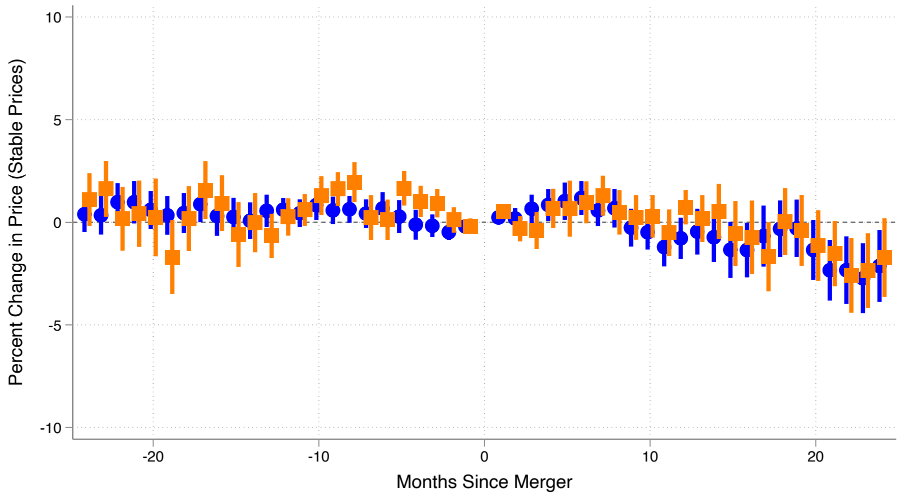

(c) Stable prices

(d) All price changes

Figure 3.5: (Continued) Timing of Price Changes, for Merging Parties (orange square) and Non-Merging Parties (blue circle)

qualitatively different. Merging parties in such mergers exhibit a decrease until about six months after the merger, after which the price stabilizes on average. We then see a further decrease by another 3–5 pp on average in the subsequent year. The price changes for non-merging parties largely mirror the prices changes for merging parties in this case. Finally, for mergers where prices are relatively stable, we do not see much movement overall. However, there is some evidence that prices decrease by 1–2 pp starting about 18 months after the merger.

One takeaway from these results is that firms do not seem to wait to exercise new market power: price increases, when they happen, happen quickly after a merger. This observation also would call into question an alternate explanation for price increases—that they reflect quality improvements rather than exercising market power—as it is unlikely that quality can be improved so quickly after a merger. On the other hand, price decreases materialize over a longer time horizon, and we find evidence that all classes of mergers exhibit an average price decrease of around 2.5% about 18–24 months after the merger. This result is consistent with synergies taking a year or more to come into play.

Panel (d) of Figure 3.5 shows the timing plot for all mergers aggregated. The salient pattern in this figure is that overall, we see a drop in prices about 12–24 months after merger completion. In fact, the average price change two years after merger completion is around -2.5%. A natural explanation in our minds is the realization of synergies on this time frame. In Section 3.5.3, we investigate more direct evidence for synergies.

3.5.2 Anticipatory Price Changes

A small set of empirical studies of airline mergers (Borenstein, 1990; Kim and Singal, 1993) and banking (Prager and Hannan, 1998) have documented evidence of anticipatory price increases, i.e., changes in behavior prior to the completion date of a proposed merger. Weinberg (2008) provides a review of these findings and notes that it is difficult to rationalize these finding with standard models of price competition, without appealing to managerial incentives.[12] Whinston (2007, p. 2427) offers a more direct explanation: that managing teams

[12]Weinberg (2008) provides one rationalization through the lens of managerial incentives and switching costs. If managers expect to get fired after a merger during restructuring, they may not have an incentive to increase market share by maintaining low prices. While an interesting rationalization for a difficult-to-explain phenomenon, it might not be broadly applicable, nor is it the only role managerial incentives may play in mergers.

spend time together before a merger and can discuss strategy. Are such anticipatory price changes common among a broader set of mergers? To answer this question, we estimate an adjustment to our main empirical specification in (3.1): we add indicators, separately for merging and non-merging parties, for the period between the merger's announcement and its completion.

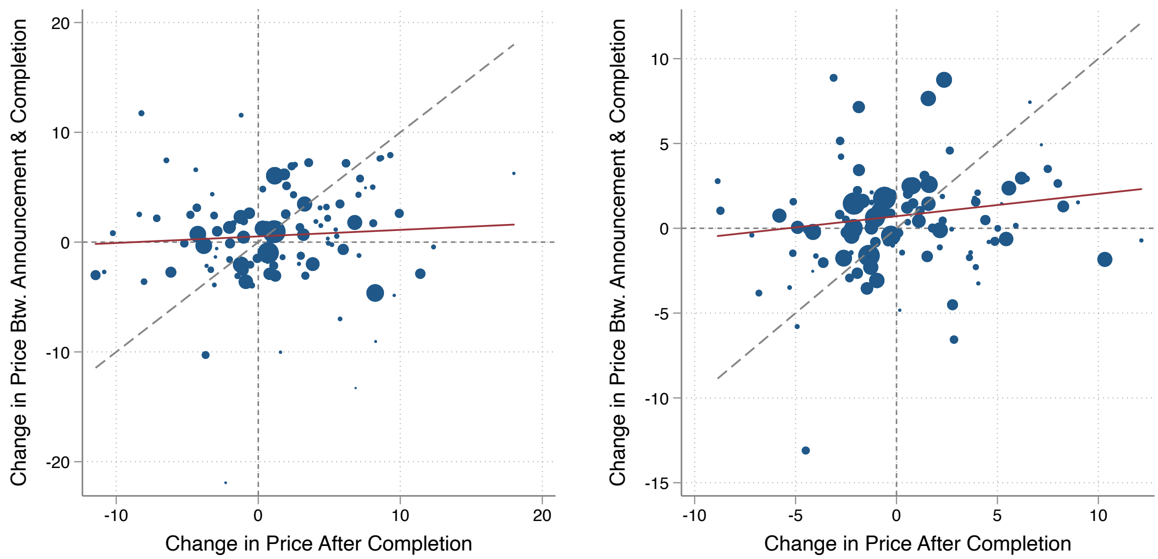

(a) Merging parties

(b) Non-merging parties

Figure 3.6: Price Changes for Merging and Non-Merging Parties Between Announcement and Completion and After Completion, as Estimated by an Adjusted Version of (3.1)

Note: The size of each point corresponds to the inverse variance of the parameter estimate for the price change after completion.

Figure 3.6 displays a scatter plot of the change in price after completion versus the change in price between the merger's announcement and its completion. We estimate both changes relative to the pre-announcement period using an adjustment to (3.1). In particular, we add a dummy for the period between the announcement date and effective date and interact it with merging and non-merging party indicators. The size of each point corresponds to the inverse variance of the post-merger period estimate.

The plots presented in Figure 3.6 do not provide evidence of uniform, anticipatory behavior across mergers. The average price change realized between the merger's announcement and execution is positive and small in magnitude for a large range of post-completion price effects. The width of the estimated distribution of anticipatory price changes is also approximately half the size of that of price changes after the merger. Moreover, we do not see any systematic patterns: it is not the case that mergers with large estimated price changes typically have large anticipatory price changes. Overall, we do not find any systematic evidence for anticipatory price changes—or for the stories of changing managerial incentives proposed by Weinberg (2008) or to the pre-merger strategy discussions referenced by Whinston (2007)—but some of the large effects estimated warrant further analysis.

3.5.3 Synergies

The decrease in prices more than a year after the merger, presented in Section 3.5.1, suggests the presence of cost synergies. However, this analysis alone does not shed light on the form they take. Empirical evidence on the magnitude and form of synergies in horizontal mergers is limited and focuses on a small number of industries.[13] Identifying the sources of synergies has important implications for the agencies' review of potential mergers, as synergies are known to be difficult to evaluate ex ante (Whinston, 2007) but frequently factor into agencies' decisions. In this section, we analyze one natural source of synergies in our setting: merger-induced changes in the geography of the merged firm's production facilities. We explore whether products sold in markets that experienced a decline in the distance to the nearest production facility faced muted price increases or larger price declines than others.

[13]For a summary, see Whinston (2007), who notes that some studies have examined "price or efficiency effects in actual mergers (none look at both). This is clearly an area that could use more research." Berger and Humphrey (1992) and Peristiani (1997) study the banking industry and do not find strong evidence of efficiencies realized in the post-merger period. Pesendorfer (2003) finds an increase in firms' efficiency in the paper industry following a set of mergers in the 1980s. In recent years, such studies have remained sparse: Ashenfelter et al. (2015) provide some evidence of efficiencies in the Miller-Coors merger while Craig et al. (2021) find no such evidence in the context of hospital mergers.

	N	Mean	S.D.	25th Pct.	Median	75th Pct.
Distance Change Definition						
< 0 miles	74	0.07	10.67	-0.97	0.38	4.89
		(2.28)		(1.37)	(2.30)	(1.55)
< -50 miles	65	-3.04	12.53	-3.57	-0.80	0.58
		(2.53)		(10.30)	(1.17)	(2.97)
< -100 miles	65	-3.05	12.65	-5.82	-0.73	0.54
		(2.58)		(10.13)	(1.22)	(2.97)
< -128 miles	63	-2.24	13.51	-1.34	-0.68	1.01
		(2.93)		(10.13)	(0.58)	(4.84)
< -200 miles	61	-2.23	14.52	-0.99	-0.61	0.95
		(3.33)		(10.74)	(0.59)	(6.51)

Table 3.5: Incremental Price Effect in Markets with Declines in Distance to the Nearest Production Facility

Note: For each distance-change classification, we display the distribution of β_4 in (3.3), weighed using the inverse variance of the estimate.

Formally, we estimate an extension of (3.1) that includes an interaction of the indicator of the post-merger period for merging firms with an indicator of the merger-induced distance change falling below a given threshold. The regression takes the form

$$\begin{aligned} \log p_{idt} = \beta_0 &+ \beta_1 \cdot \mathbb{1}[\text{Merging Party}]_i + \beta_2 \cdot \mathbb{1}[\text{Post-Merger}]_t \cdot \mathbb{1}[\text{Merging Party}]_i \\ &+ \beta_3 \cdot \mathbb{1}[\text{Post-Merger}]_t \cdot \mathbb{1}[\text{Non-Merging Party}]_i \\ &+ \beta_4 \cdot \mathbb{1}[\text{Post-Merger}]_t \cdot \mathbb{1}[\text{Merging Party}]_i \cdot \mathbb{1}[\Delta\text{Distance} < \bar{D}]_{id} \\ &+ \xi_{id} + \xi_{m(t)} + \text{Controls}_{idt} + \epsilon_{idt}, \end{aligned} \tag{3.3}$$

where $\bar{D}$ is the specified distance threshold (in miles), and $\bar{D} \in \{0, -50, -100, -128, -200\}$.[14] In this specification, β_4 is the coefficient of interest, which measures the additional price effect on products sold in markets that experienced sufficiently large declines in the distance to the nearest production facility.

[14]The first quartile of the distribution of distance changes across all mergers is -128 miles. We specify this as one value of $\bar{D}$ and examine robustness to a range of others.

Table 3.5 displays the distribution of the incremental price effect in markets that experience declines in distance to the nearest production facility, i.e., β_4 of (3.3). We compute the distributions separately for different definitions of distance changes and weigh using the inverse variance of the estimate. We observe lower price changes, on average, in markets characterized by a decline in the distance to the nearest production facility of more than 50 miles. This result is suggestive of the realization of cost synergies through changes in the firms' distribution networks. It is important to note, however, that there is substantial heterogeneity across mergers in both the direction and magnitude of the effect. This variation could stem, for example, from differences in transportation costs across product types. Miller and Weinberg (2017) note the importance of transportation costs in beer distribution, but freight costs vary substantially across products (Behrens et al., 2018).

The results in this section highlight one potential way in which synergies could manifest themselves in mergers of consumer packaged goods firms. Following the completion of a merger, firms have the option to reallocate production across their new set of facilities. We find evidence that the resulting reallocation is passed through to prices, on average. However, as explained above, this is not the case across all mergers. Together, these results point to future avenues of research to determine: (1) which subset of firms and products are most likely to benefit from distributional synergies, and (2) whether changes to the geography of production allow firms to better cater to region-specific tastes. Work is in progress to address these two points.

3.5.4 Correlation With Market Structure

Our final analysis studies the market-power channel more directly, as we examine the correlation between the observed price changes and standard measures of the level and change in market concentration (i.e., post-merger HHI and ΔHHI). Why focus on these metrics? When evaluating a merger, antitrust agencies typically weigh the potential to realize cost synergies like those described in Section 3.5.3 with changes in market structure. The latter piece is codified in the 2010 Horizontal Merger Guidelines, which detail presumptions by which the agencies identify mergers likely to lead to consumer harm.[15] Therefore, standard measures of market structure frequently enter into the agencies' decisions regarding whether to approve a merger or require an enforcement action.

The theoretical basis of these presumptions has been the focus of recent work in the academic literature. Theoretical results show a relationship between ΔHHI and the price and consumer surplus effects of mergers (Nocke and Schutz, 2018; Nocke and Whinston, 2021). On the other hand, Nocke and Whinston (2021) call into question the basis of basing merger screens on levels of HHI: they show that under both stylized theoretical demand systems and tractable empirical ones, the surplus effects of a merger are not related to HHI. Nevertheless, one may imagine alternate reasons that HHI would play a role in the effects of mergers: for instance, Loertscher and Marx (2021) note that competition authorities have historically used HHI as a measure of the potential for coordinated effects (but also call this into question). Ultimately, evaluating the screens is fundamentally an empirical question, and despite the recent attention paid to these presumptions, there is little evidence examining their efficacy.[16] In this section, we provide large-scale empirical evidence that the price effects of a merger correlate with the nationwide change in HHI but not with its post-merger level. These findings, however, do not hold across geographic markets within a merger.

Figure 3.7 illustrates the relationship between nationwide concentration metrics and price changes. Specifically, it presents a scatter plot of merging and non-merging party price

[15] Mergers that change HHI by more than 100 points and lead to a post-merger HHI between 1,500 and 2,500 or change HHI by between 100 and 200 points and lead to a post-merger HHI above 2,500 "raise significant competitive concerns and often warrant scrutiny." Mergers that change HHI by more than 200 points and lead to a post-merger HHI above 2,500 are "presumed to be likely to enhance market power." U.S. Department of Justice and the Federal Trade Commission (2010): "Horizontal Merger Guidelines," available at https://www.justice.gov/atr/horizontal-merger-guidelines-08192010#1.

[16] Majerovitz and Yu (2021) find statistically insignificant differences in the merger effects on price and revenue across values of ΔHHI, but their sample contains mergers with small changes in HHI, on average (mean of 4.14). Our sample contains more variation in ΔHHI.

changes across HHI and ΔHHI. Recall that the structural presumptions rely on cutoffs in both HHI and ΔHHI: mergers raise competitive concerns if they are large on both dimensions. The different colors and shapes represent the quintile of the distribution of price changes for merging parties. Movement in the color from blue to red indicates higher average price changes for the given set of mergers. Panel (a) presents the plot for all mergers, while Panel (b) restricts to mergers with $\Delta\text{HHI} < 400$ and post-merger $\text{HHI} < 7000$ to remove the influence of outliers. We present the analogous plots for non-merging parties in Figure C.2, as well as heat-map versions in Figure C.3. Generally, price changes are larger for mergers with larger ΔHHI, especially for low HHI: the red dots are more concentrated above low values of ΔHHI. However, we do not see a similar pattern for HHI.

Figure 3.8 illustrates another cut of this relationship: the horizontal axis in each panel plots either HHI or ΔHHI, and the vertical axis plots the average price change of all mergers that have a value at least as large. Panel (a) shows the graphs as a function of ΔHHI, and we generally see price changes increase with ΔHHI. Panel (b) shows these price changes as a function of HHI, and we see that they are fairly flat. In aggregate, these results are consistent with the theoretical work of Nocke and Whinston (2021) and suggest that ΔHHI (at the national level) correlates with the parties' ability to exploit market power afforded them by a merger.

Panel (c) cuts the relationship shown in Panel (a) by HHI (splitting on an HHI of 2,500). We see that the slope is larger for mergers with low HHI. One potential explanation for the difference in these relationships is a form of selection. Mergers with high post-merger HHI and ΔHHI likely face greater scrutiny from the antitrust agencies and therefore may require a higher burden of proof that such deals would produce cost synergies. On the other hand, we see no clear pattern for HHI, even when we split Panel (b) at ΔHHI equal to 200 (Panel (d)).

We have thus far focused on variation in price effects across mergers. How do price effects

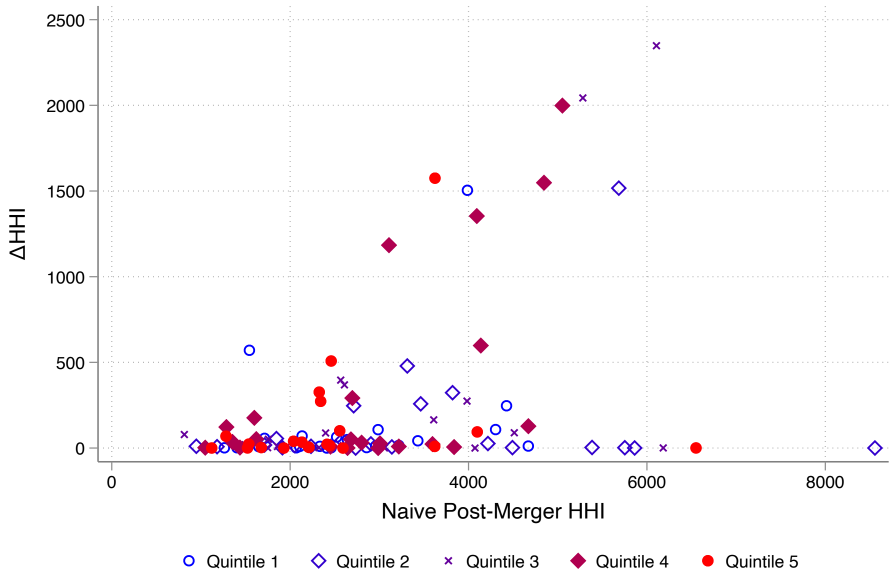

(a) Merging, Unrestricted

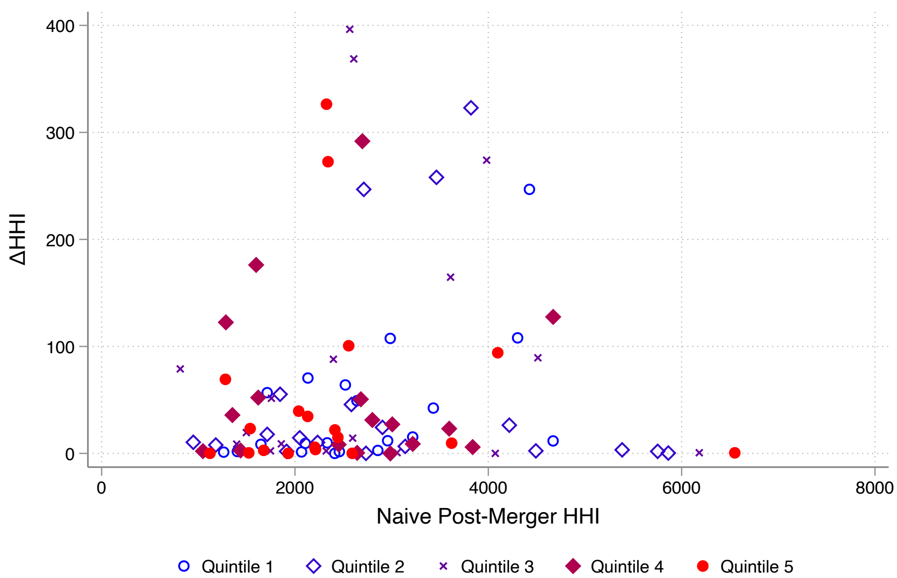

(b) Merging, Restricted

Figure 3.7: Scatter Plots of Nationwide ΔHHI Against Nationwide Post-Merger HHI for Mergers in Our Sample

Note: Panel (a) displays the results for all mergers, while Panel (b) restricts to mergers for which $\Delta\text{HHI} < 400$ and post-merger $\text{HHI} < 7,000$. Colors/shapes represent the quintile of the price change for merging parties, with solid (red) colors representing larger price changes. Note that $\Delta\text{HHI} < 0$ corresponds to divestitures being required in the merger.

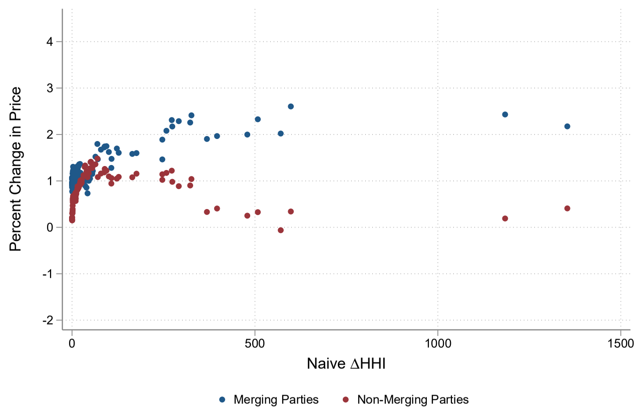

(a) Naive ΔHHI

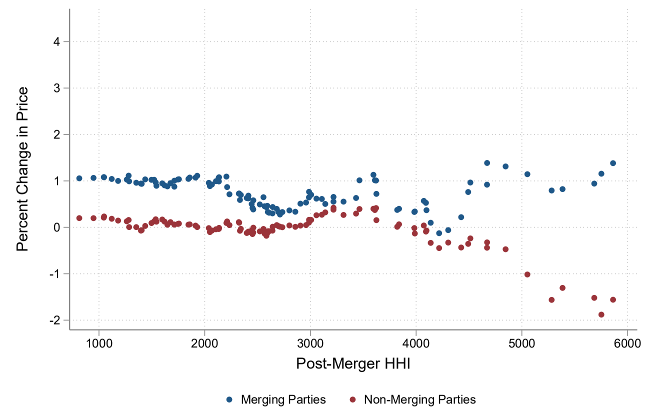

(b) Post-merger HHI

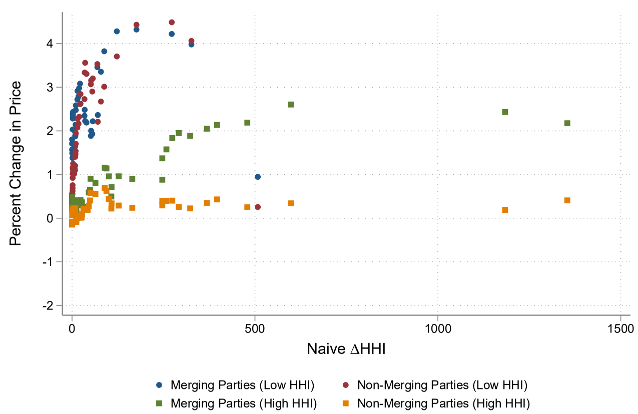

(c) Naive ΔHHI, by HHI $\gtrless$ 2500

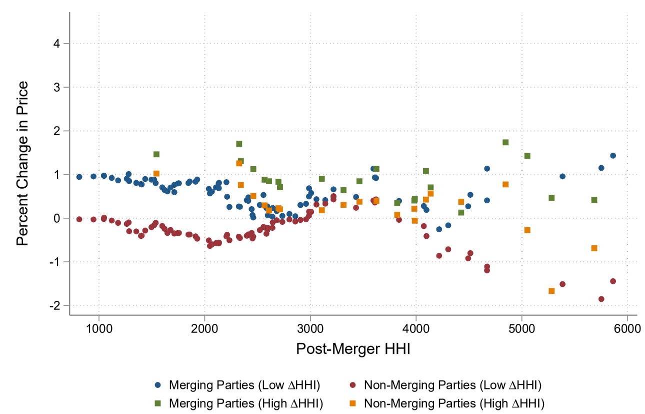

(d) Post-merger HHI, by ΔHHI $\gtrless$ 200

Figure 3.8: Average Price Changes of a Merger with At Least a Given Level of ΔHHI ((a) and (c)) or Post-Merger HHI ((b) and (d))

Note: Panels (c) and (d) split the sample by levels of HHI or ΔHHI.

differ within-merger across geographic markets (i.e., DMAs), which also exhibit heterogeneity in HHI and ΔHHI due to the merger? We estimate an adjustment to (3.1), where we interact the post-merger dummy for merging and non-merging parties with indicators of whether a market falls in given ΔHHI or post-merger HHI category.[17] In Figure 3.9, we plot the mean price change, across mergers, in each category, weighed by the inverse variance of the parameter estimate. Orange squares correspond to merging parties, and blue circles correspond to non-merging parties. The solid lines indicate the 90% confidence interval of that mean.

While results are somewhat noisy, we see no clear evidence of a gradient in price effects by ΔHHI or by HHI of the market. We find some evidence of an inverted-U shape, at least for merging parties: we see no price effects for markets with especially low or especially high HHI or ΔHHI, but prices increase in markets with intermediate levels of these variables. One explanation for these effects may again be selection. Given that antitrust agencies have the ability to require divestitures only in a subset of geographies (although this power is exercised only once in our dataset), the fact that agencies allowed mergers to go through even though certain markets had large increases in concentration may be indicative of them anticipating large synergies. Second, the set of mergers that contribute to each estimate is different, as not all mergers have markets in each bin. This induces another source of selection: a low estimate of the price effect for high HHI could be explained by mergers with markets large HHI also having low overall ΔHHI.[18] In Section 3.6, we develop a model that accounts for selection into merger approval.

It is important to note that we do not assign a causal interpretation to these findings. The

[17] We group ΔHHI in buckets of size 25 from 0 to 200. For post-merger HHI, we use the following cutoffs: (375, 750, 1,125, 1,500, 1,750, 2,000, 2,250, and 2,500).

[18] Additionally, while we do not find strong evidence of this in Section 3.4.2, some firms may still use coarse pricing rules. If they set prices regionally or nationally and thus do not respond to DMA-level changes in concentration due to the merger, the resulting spillovers across geographies could confound the results presented in Figure 3.9.

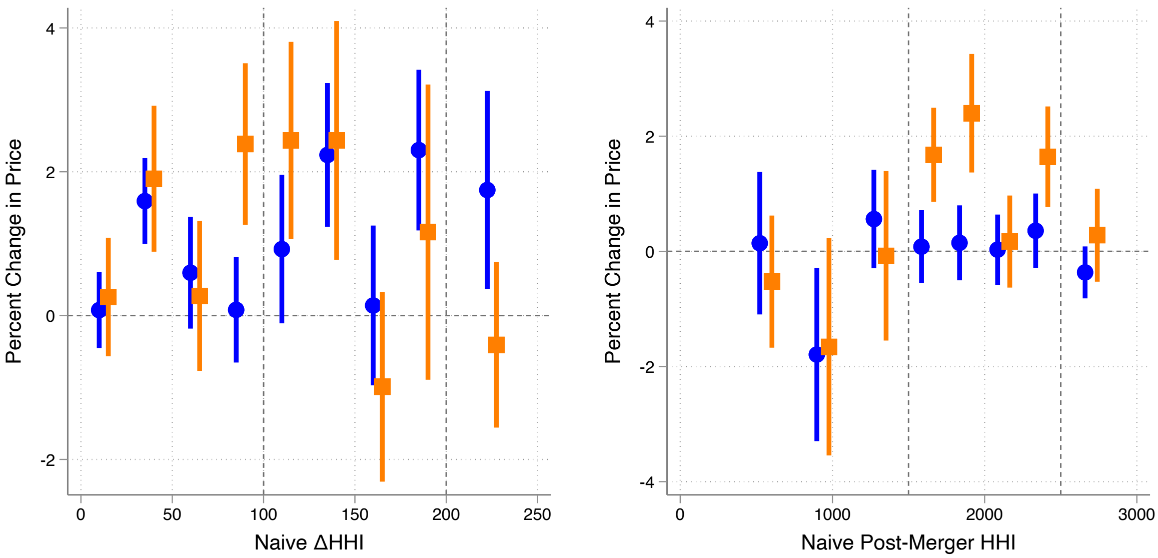

(a) ΔHHI

(b) Post-Merger HHI

Figure 3.9: Prices Changes Post-Merger, by Bins of Naive ΔHHI and HHI, for Merging Parties (orange square) and Non-Merging Parties (blue circle)

Note: The marker indicates the mean price change, weighed by the inverse variance of the parameter estimate, for geographic markets in that bin across all mergers, and the thick line is the 90% confidence interval of that mean.

industrial organization literature has long cautioned against regressions of price or quantity on measures of market concentration (see, e.g., Demsetz (1973), Schmalensee (2012), and Berry et al. (2019)). Price, quantity, and HHI are all equilibrium outcomes, and it is difficult to define an ideal experiment that varies HHI. However, the agencies include HHI and ΔHHI in their presumptions of consumer harm, and it is plausible that they relate to the merged entity's ability to exercise market power. Therefore, it is still useful to provide large-scale evidence of whether these measures correlate with price and quantity changes.

3.6 How Stringent is US Antitrust Enforcement?

Analyzing the universe of large mergers avoids selection into publication: the effects we document are necessarily representative of mergers in US retail above the chosen deal size. This does not mean, however, that they are representative of the effects of all possible mergers—even between firms similar to the ones in our dataset. The effects that we have documented thus far are mediated through two important channels of selection. First, the merging parties have to propose the merger, meaning it must be profitable. Second, the enforcement agencies must approve it (or lose a challenge in court), meaning that there should be an expectation that it would not be too harmful to consumer surplus.

In this section, we delve into the second source of selection mentioned above: what types of mergers are the agencies willing to approve? As noted in Carlton (2009), noting that the mean price change is small—or even that the a large set of mergers have substantial price decreases—does not suggest that the enforcement agencies are strict, challenging mergers that have large price increases (or perhaps dissuading them from being proposed at all). Carlton (2009) notes that the stringency of an enforcement agency can be measured by the "marginal" merger: in a simple world in which agencies could unilaterally approve or reject a merger, and where merger effects could be perfectly predicted, this would be the largest price effect observed among consummated mergers. However, in a more realistic setup, the existence of consummated mergers with large price changes could also be indicative of uncertainty.

Conceptually, we can model the agencies as choosing to "challenge" mergers that they believe to be sufficiently anticompetitive—that they expect will lead to significant price increases. Generally, a challenge could be one of many actions, e.g., a motion to block the merger or a proposal for a remedy. In our setting, a "challenge" always involves proposing a remedy: the agencies did not (unsuccesfully) try to block any of the mergers in our dataset, but they did propose remedies in four of the deals in our dataset.[19] The agency has some information X_i about the merger i, which gives it a prior for the true price change p_i^* of the merger. They also get a noisy signal of the price change of the merger based on further due diligence conducted and then decide whether to challenge the merger.

In a general model, the agency effectively has a probability $\lambda(p_i^*, X_i)$ of challenging a merger where the true price change is p_i^* and observable characteristics are X_i. The randomness in this decision (when viewed from the perspective of the econometrician) could come from two sources: (i) noise in due diligence or (ii) characteristics that are unobserved to the econometrician but used in the agencies' decision. In this section, we estimate an empirical version of the model by imposing the first interpretation. Appendix C.2 discusses the connection to the second interpretation.

We provide a simple empirical implementation of this model of antitrust enforcement. The true price impact of merger i, with observed characteristics X_i, is p_i^*. Agencies have a prior that this price impact is drawn from $N(\mu_{p^*}(X_i), \sigma_{p^*}^2)$, based on the price impacts of mergers that were similar on observables. From further investigation of the proposed merger

[19]SDC Platinum also lists deals that were withdrawn, although they do not systematically list whether the reason was antitrust scrutiny. We are in the process of identifying whether the reason for withdrawal was antitrust scrutiny; we will include information for deals for which we believe this to be the case and treat them as deals the agencies "blocked."

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
True Price Change								
Constant	0.44	0.37	0.92	0.92	1.16	1.11	1.63	1.63
	[-0.31, 1.20]	[-0.37, 1.13]	[-0.49, 2.33]	[-0.49, 2.33]	[0.13, 2.21]	[0.09, 2.16]	[-0.31, 3.57]	[-0.30, 3.57]
HHI			-2.02				-2.14	
			[-6.48, 2.47]				[-8.29, 4.00]	
DHHI			3.25				4.09	
			[-3.66, 10.29]				[-5.24, 13.60]	
Mean HHI				-2.03				-2.16
				[-6.48, 2.46]				[-8.30, 3.98]
Mean DHHI				3.18				4.00
				[-3.73, 10.21]				[-5.35, 13.51]
σ_{p^*}	3.87	3.81	3.84	3.84	5.30	5.25	5.23	5.23
	[3.35, 4.48]	[3.32, 4.40]	[3.33, 4.45]	[3.33, 4.44]	[4.57, 6.15]	[4.53, 6.09]	[4.53, 6.06]	[4.52, 6.05]
Price Threshold								
Constant	12.75	20.02	10.48	10.06	12.16	12.76	11.18	10.81
	[7.76, 30.08]	[6.99, 66.10]	[7.30, 16.13]	[6.84, 15.54]	[8.80, 20.24]	[7.54, 32.05]	[8.15, 16.16]	[7.79, 15.68]
Log(Sales)	-1.63		-1.41	-1.36	-1.60		-1.49	-1.44
	[-4.76, 0.21]		[-3.29, 0.00]	[-3.28, 0.11]	[-3.75, -0.11]		[-3.21, -0.21]	[-3.19, -0.12]
Log(Merg Sales)		-5.53				-3.17		
		[-22.11, -0.69]				[-9.95, -0.82]		
HHI			1.69				2.39	
			[-5.69, 9.87]				[-4.33, 9.19]	
DHHI			-22.47				-19.10	
			[-40.51, -12.20]				[-35.69, -9.86]	
Mean HHI				2.75				3.37
				[-4.48, 11.24]				[-3.09, 10.37]
Mean DHHI				-22.59				-19.20
				[-40.73, -12.32]				[-35.91, -9.81]
σ_ϵ	6.01	13.22	4.01	4.03	4.28	6.60	3.42	3.43
	[2.50, 16.83]	[3.05, 46.70]	[2.18, 7.24]	[2.20, 7.26]	[1.94, 10.03]	[2.16, 21.34]	[1.77, 6.37]	[1.76, 6.41]
Eff. Threshold	4.56	3.69	5.61	5.63	8.20	8.18	8.75	8.83
	[1.74, 7.31]	[0.73, 7.72]	[3.25, 8.10]	[3.27, 8.14]	[4.85, 10.77]	[3.24, 12.16]	[6.16, 11.11]	[6.21, 11.21]
Price	Overall	Overall	Overall	Overall	Merging	Merging	Merging	Merging

Table 3.6: Posterior Means and 95% Credible Intervals of the Model of Selection into Enforcement

Note: Total pre-merger sales ("Sales") and pre-merger sales of the merging party ("Merg Sales") are in billions of dollars. HHI and (naive) DHHI are calculated at the national level, and Mean HHI and Mean DHHI are averaged over markets, in the pre-merger period. Columns (1)–(4) use average price changes as measured in a regression of a post-merger dummy on UPC-DMA fixed effects, month fixed effects, and a DMA-specific time trend. Columns (5)–(8) use the price change of merging parties, as measured in the baseline.

at hand, they also observe a noisy signal $p_i = p_i^* + \epsilon_i$, with $\epsilon_i \sim N(0, \sigma_\epsilon^2)$. Agencies will ask for remedies if their signal of the price impact is larger than some threshold, i.e., if $p_i \geq \bar{p}(X_i)$.

In our setting, we observe two important pieces of information for each merger: (i) an estimate of the price change due to the merger (overall, and of merging parties in particular), and (ii) whether the agency proposed a remedy for it. We also observe some characteristics X_i about each merger i, such as the market size and market structure. We assume that our estimate of the price change is $\hat{p}_i \sim N(p_i^*, \sigma_i^2)$, where σ_i is the standard error of our estimate of the price effect. We do so only for mergers that were not challenged: otherwise we only observe an estimate of the price effect after the remedy, and we wish to avoid taking a stance on how this price change would compare with the one without a remedy.

To gain intuition for how the data informs the parameters of this model, suppose that true price changes are observed for consummated mergers. Suppose in addition that there is a merger-specific property Z_i that affects the agencies' threshold $\bar{p}(\cdot)$ but not the prior distribution of expected price changes. Consider the case where $\sigma_\epsilon = 0$: the agency can also perfectly predict the price change. When Z_i is such that the agency is especially permissive, we observe the distribution of price changes unfiltered through the agencies' actions: this is a good estimate of the prior distribution itself and thus of $(\mu_{p^*}, \sigma_{p^*})$. As Z_i changes and the agency grows less permissive, the maximum observed price change among all consummated mergers informs $\bar{p}$. In the more realistic case with noise in the agencies' assessments ($\sigma_\epsilon > 0$), one can compare the observed price change distribution for a stringent Z_i with one for a permissive Z_i: if the distributions are similar, then there is still a fair amount of residual noise in the agencies' assessment.

We parameterize $\mu_{p^*}(\cdot)$ and $\bar{p}(\cdot)$ as linear combinations of merger characteristics. In the most general specification, we let both depend on HHI and ΔHHI, either at the merger level or an average over the values for geographic markets. We assume that the characteristic Z_i that enters the price threshold but is excluded from the price change distribution is total sales in the market. The rationale behind this choice is that antitrust agencies may be more likely to scrutinize mergers in larger markets even conditional on concentration. We estimate the model using a computational Bayesian approach of Gibbs sampling with data augmentation. Table 3.6 shows results.

The first set of columns use the average price change of all parties (merging and non-merging) as the price change of interest, estimated using a specification with UPC-DMA fixed effects, month fixed effects, and DMA-specific time trends. Column (1) shows a specification where μ_{p^*} is parameterized with a constant, and the price threshold $\bar{p}$ depends on the total pre-merger sales in the market. Consistent with our intuition, we find a negative point estimate: a 10% increase in market size reduces the threshold by 0.16 pp. Column (2) uses the total pre-merger sales by merging parties in the market, and we again find a negative relation, although the magnitude is about three times as large. To interpret the magnitude of the threshold, we compute the expected price change for a merger whose signal p_i equals the threshold $\bar{p}(X_i)$; this is the posterior mean $(\bar{p}(X_i) \cdot \sigma_\epsilon^{-2} + \mu_{p^*}(X_i) \cdot \sigma_{p^*}^{-2})/(\sigma_\epsilon^{-2} + \sigma_{p^*}^{-2})$. The effective threshold, reported in the final row of Table 3.6 is the average of this quantity over all mergers in the sample. In Columns (1) and (2), we find thresholds of 4.6% and 3.7%, respectively; overall, enforcement agencies propose remedies for any merger whose expected price increase exceeds this threshold. The credible intervals for these estimates are somewhat large, but they do not contain 0.

Columns (3) and (4) add measures of pre-merger market structure to both the mean of the true price distribution and the threshold—either the nationwide pre-merger HHI and (naive) ΔHHI or the average of these measures at the market-level. We do not document any significant relationship between the measures of market structure and the prior of the price distribution, but we consistently find a significant negative relationship between ΔHHI and the price threshold. A ΔHHI that is 100 points larger corresponds to a reduction in the threshold of about 0.2 pp—consistent with the observation that the agencies are more likely to scrutinize mergers with larger changes in concentration. In these specifications, we estimate a slightly larger effective threshold of about 5.6%.

Finally, Columns (5)–(8) repeat the specifications in Columns (1)–(4) but use the price change of merging parties as the effect of interest, estimated from the baseline. Results are qualitatively similar, but we estimate a larger effective threshold of 8.2–8.8% depending on the specification.

An interpretation of this result is that the "marginal" merger would have a price effect in the range of 3.7–4.6% overall, or with a price increase of 8.2–8.8% for merging parties. Kwoka (2014, p. 86) argues that one interpretation of the selection bias in published studies is that these studies are more likely to be of such marginal mergers: these are the deals that garnered press attention partly because of agency scrutiny. It is thus noteworthy that he arrives at a quantitatively similar conclusion, with mean price changes of mergers around 7.2% (Table 7.2 in Kwoka (2014)).[20]

3.7 Conclusion

In this paper, we conduct a large-scale retrospective of all sufficiently large mergers in US retail. Over a number of specifications, we find a relatively modest average effect of 0.5% on price and a decrease in quantity of around 3%. These means mask substantial heterogeneity in outcomes, with an interquartile range of almost 5% for price changes. These average effects are nevertheless considerably smaller than those reported in previous meta-analyses of

[20] Another interpretation places some caveats on the threshold interpretation. A general model is that the probability of challenging a merger is a function of the true price change. Appendix C.2 shows that the probability function is identified from data, and this threshold is simply one way to interpret this function. However, the decision to approve or reject need not be solely due to a noisy estimate of the price change: the agencies may well have more information about the merger, on which it bases this decision. We imagine this is realistic, not just for institutional reasons (such as the fact that agencies conduct diligence on proposed mergers and receive more information from parties in the HSR filings) but also because we estimate σ_ϵ to be roughly the same magnitude as σ_{p^*} in most specifications. If other factors are part of this ϵ, we can roughly interpret the threshold reported in this section as an average of thresholds conditional on merger-specific information unobserved to us.

merger retrospectives; the difference is that our sample consists of inframarginal mergers—all mergers that were proposed and allowed by the enforcement agencies—rather than consistent of a selected sample of mergers that may have garnered particular interest from the agencies, the press, or researchers. We show that this mean is by itself not informative of the agencies' stringency. We use a simple selection model of the agencies' enforcement decision to estimate that agencies tend to suggest remedies for mergers they expect will increase average prices by 3.7–5.6%.

References

Adams, B. and K. R. Williams (2019). Zone Pricing in Retail Oligopoly. *American Economic Journal: Microeconomics 11*(1), 124–56.

Allen, J., R. Clark, and J. Houde (2019). Search Frictions and Market Power in Negotiated Price Markets. *Journal of Political Economy 127*(4), 1550–1598.

Ashenfelter, O. C. and D. S. Hosken (2010). The Effect of Mergers on Consumer Prices: Evidence from Five Mergers on the Enforcement Margin. *Journal of Law and Economics 53*(3), 417–466.

Ashenfelter, O. C., D. S. Hosken, and M. C. Weinberg (2013). The Price Effects of a Large Merger of Manufacturers: A Case Study of Maytag-Whirlpool. *American Economic Journal: Economic Policy 5*(1), 239–61.

Ashenfelter, O. C., D. S. Hosken, and M. C. Weinberg (2014). Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consummated Mergers. *Journal of Law and Economics 57*(S3), S67–S100.

Ashenfelter, O. C., D. S. Hosken, and M. C. Weinberg (2015). Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry. *RAND Journal of Economics 46*(2), 328–361.

Asker, J. and V. Nocke (2021). Collusion, Mergers, and Related Antitrust Issues. In *Handbook of Industrial Organization*, Volume 5, pp. 177–279. Elsevier.

Atalay, E., A. Sorensen, C. Sullivan, and W. Zhu (2020). Post-Merger Product Repositioning: An Empirical Analysis. *Working Paper*.

Autor, D., D. Cho, L. D. Crane, B. Lutz, J. Montes, W. B. Peterman, D. Ratner, D. Villar, and A. Yildirmaz (2020). An Evaluation of the Paycheck Protection Program Using Administrative Payroll Microdata. *Working Paper*.

Avramidis, P., K. Serfes, and K. Wu (2021). The Role of Regulation and Competition in Credit Allocation: Evidence from Small Business Lending. *Working Paper*.

Bachas, N., O. Kim, and C. Yannelis (2021). Loan Guarantees and Credit Supply. *Journal of Financial Economics 139*(3), 872–894.

Balyuk, T., N. Prabhala, and M. Puri (2020). Indirect Costs of Government Aid and Intermediary Supply Effects: Lessons From the Paycheck Protection Program. *Working Paper*.

Barraza, S., M. Rossi, and T. J. Yeager (2020). The Short-Term Effect of the Paycheck Protection Program on Unemployment. *Working Paper*.

Bartik, A. W., Z. B. Cullen, E. L. Glaeser, M. Luca, C. T. Stanton, A. Sunderam, and W. Bartik (2020). The Targeting and Impact of Paycheck Protection Program Loans to Small Businesses. *Working Paper*.

Behrens, K., W. M. Brown, and T. Bougna (2018). The World Is Not Yet Flat: Transport Costs Matter! *The Review of Economics and Statistics 100*(4), 712–724.

Berger, A. N. and D. B. Humphrey (1992). Megamergers in Banking and the Use of Cost Efficiency as an Antitrust Defense. *The Antitrust Bulletin 37*(3), 541–600.

Berry, S., M. Gaynor, and F. Scott Morton (2019). Do Increasing Markups Matter? Lessons from Empirical Industrial Organization. *Journal of Economic Perspectives 33*(3), 44–68.

Berry, S. T. and J. Waldfogel (2001). Do Mergers Increase Product Variety? Evidence from Radio Broadcasting. *Quarterly Journal of Economics 116*(3), 1009–1025.

Borenstein, S. (1990). Airline Mergers, Airport Dominance, and Market Power. *The American Economic Review 80*(2), 400–404.

Brown, J. and J. Earle (2017). Finance and Growth at the Firm Level: Evidence from SBA Loans. *Journal of Finance 72*(3), 1039–1080.

Callaway, B., A. Goodman-Bacon, and P. Sant'Anna (2021). Difference-in-Differences with a Continuous Treatment. *Working Paper*.

Carlton, D. (2009). The Need to Measure the Effect of Merger Policy and How to Do It. *Competition Policy Internation 5*(1), Article 6.

Chernenko, S. and D. S. Scharfstein (2021). Racial Disparities in the Paycheck Protection Program. *Working Paper*.

Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, and The Oppourtnity Insights Team (2020). How Did Covid-19 and Stabilization Policies Affect Spending and Employment? a New Real-Time Economic Tracker Based on Private Sector Data. *Working Paper*.

Chiappori, P. and B. Salanie (2000). Testing for Asymmetric Information in Insurance Markets. *Journal of Political Economy 108*(1), 56–78.

Choi, B. and H. Lee (2019). Heterogeneous Elasticities of Bank Loan Supply: Evidence from SBA Loan Guarantee Expansion. *Working Paper*.

Congressional Research Service (2019a, 08). Small Business Administration 7(a) Loan Guaranty Program. CRS Report.

Congressional Research Service (2019b, 03). Small Business Administration: A Primer on Programs and Funding. CRS Report.

Cororaton, A. and S. Rosen (2021). Public Firm Borrowers of the US Paycheck Protection Program. *The Review of Corporate Finance Studies 10*(4), 641–693.

Cox, N., E. Liu, and D. Morrison (2021). Market Power in Small Business Lending: A Two-Dimensional Bunching Approach. *Working Paper*.

Craig, S. V., M. Grennan, and A. Swanson (2021). Mergers and Marginal Costs: New Evidence on Hospital Buyer Power. *RAND Journal of Economics 52*(1), 151–178.

Crawford, G., N. Pavanini, and F. Schivardi (2018). Asymmetric Information and Imperfect Competition in Lending Markets. *American Economic Review 108*(7), 1659–1701.

Cuesta, J. and A. Sepulveda (2021). Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access. *Working Paper*.

Dafny, L., M. Duggan, and S. Ramanarayanan (2012). Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry. *American Economic Review 102*(2), 1161–85.

Das, S. (2019). Effect of Merger on Market Price and Product Quality: American and US Airways. *Review of Industrial Organization 55*(3), 339–374.

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

Demsetz, H. (1973). Industry Structure, Market Rivalry, and Public Policy. *Journal of Law and Economics 16*(1), 1–9.

Deneckere, R. and C. Davidson (1985). Incentives to Form Coalitions with Bertrand Competition. *RAND Journal of economics*, 473–486.

Denes, M., R. Duchin, and J. Hackney (2021). Does Size Matter? The Real Effects of Subsidizing Small Firms. *Working Paper*.

DerSimonian, R. and N. Laird (1986). Meta-Analysis in Clinical Trials. *Controlled Clinical Trials 7*(3), 177–188.

Dilger, R. (2016). Small Business Administration 7(a) Loan Guaranty Program. *Congressional Research Service*.

Doniger, C. and B. Kay (2021). Ten Days Late and Billions of Dollars Short: The Employment Effects of Delays in Paycheck Protection Program Financing. *Working Paper*.

Draganska, M., M. Mazzeo, and K. Seim (2009). Beyond Plain Vanilla: Modeling Joint Product Assortment and Pricing Decisions. *Quantitative Marketing and Economics 7*(2), 105–146.

Edelberg, W. (2006). Risk-Based Pricing of Interest Rates for Consumer Loans. *Journal of Monetary Economics 53*(8), 2283–2298.

Einav, L., M. Jenkins, and J. Levin (2012). Contract Pricing in Consumer Credit Markets. *Econometrica 80*(4), 1387–1432.

Eliason, P. J., B. Heebsh, R. C. McDevitt, and J. W. Roberts (2020). How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry. *Quarterly Journal of Economics 135*(1), 221–267.

Erel, I. and J. Liebersohn (2020). Does Fintech Substitute for Banks? Evidence from the Paycheck Protection Program. *Working Paper*.

Fairlie, R. W. and F. M. Fossen (2021). Did the $660 Billion Paycheck Protection Program and $220 Billion Economic Injury Disaster Loan Program Get Disbursed to Minority Communities in the Early Stages of COVID-19? *Working Paper*.

Fan, Y. (2013). Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market. *American Economic Review 103*(5), 1598–1628.

Farrell, J., P. A. Pautler, and M. G. Vita (2009). Economics at the FTC: Retrospective Merger Analysis with a Focus on Hospitals. *Review of Industrial Organization 35*(4), 369.

Faulkender, M., R. Jackman, S. Miran, R. Jackman, and S. Miran (2021). The Job-Preservation Effects of Paycheck Protection Program Loans. *Working Paper* (April).

Focarelli, D. and F. Panetta (2003). Are Mergers Beneficial to Consumers? Evidence from the Market for Bank Deposits. *American Economic Review 93*(4), 1152–1172.

Garmon, C. (2017). The Accuracy of Hospital Merger Screening Methods. *RAND Journal of Economics 48*(4), 1068–1102.

Gorton, G. and G. Pennacchi (1995). Banks and Loan Sales Marketing Nonmarketable Assets. *Journal of Monetary Economics 35*(3), 389–411.

Granja, J., C. Makridis, C. Yannelis, and E. Zwick (2021). Did the Paycheck Protection Program Hit The Target? *Working Paper*.

Greene, W. (2012). *Econometric Analysis* (7 ed.). Prentice Hall.

Haas-Wilson, D. and C. Garmon (2011). Hospital Mergers and Competitive Effects: Two Retrospective Analyses. *International Journal of the Economics of Business 18*(1), 17–32.

Hauswald, R. and R. Marquez (2006). Competition and Strategic Information Acquisition in Credit Markets. *The Review of Financial Studies 19*(3), 967–1000.

Holmstrom, B. and J. Tirole (1997). Financial Intermediation, Loanable Funds, and the Real Sector. *The Quarterly Journal of Economics 112*(3), 663–691.

Howell, S. T., T. Kuchler, D. Snitkof, J. Stroebel, and J. Wong (2021). Racial Disparities in Access to Small Business Credit: Evidence from the Paycheck Protection Program. *Working Paper*.

Hubbard, G. and M. R. Strain (2020). Has the Paycheck Protection Program Succeeded? *Working Paper*.

Humphries, J. E., C. A. Neilson, and G. Ulyssea (2020). Information frictions and access to the Paycheck Protection Program. *Journal of Public Economics 190*, 104244.

Hunter, G., G. K. Leonard, and G. S. Olley (2008). Merger Retrospective Studies: A Review. *Antitrust 23*, 34.

Ioannidou, V., J. Liberti, T. Mosk, and J. Sturgess (2018). Intended and Unintended Consequences of Government Credit Guarantee Programs. In C. Mayer, S. Micossi, M. Onado, M. Pagano, and A. Polo (Eds.), *Finance and Investment: The European Case*, pp. 317–325. Oxford: Oxford University Press.

Ioannidou, V., N. Pavanini, and Y. Peng (2022). Collateral and Asymmetric Information in Lending Markets. *Journal of Financial Economics 144*(1), 93–121.

James, C., J. Lu, and Y. Sun (2021). Time is Money: Real Effects of Relationship Lending in a Crisis. *Journal of Banking and Finance 133*, 1–21.

Joaquim, G. and F. Netto (2021). Optimal Allocation of Relief Funds: The Case of the Paycheck Protection Program. *Working Paper*.

Keys, B., T. Mukherjee, A. Seru, and V. Vig (2010). Did Securitization Lead to Lax Screening? Evidence from Subprime Loans. *The Quarterly Journal of Economics 125*(1), 307–362.

Kim, E. H. and V. Singal (1993). Mergers and Market Power: Evidence from the Airline Industry. *American Economic Review*, 549–569.

Kwoka, J. (2014). *Mergers, Merger Control, and Remedies: A Retrospective Analysis of US Policy.* MIT Press.

Kwoka, J. and E. Shumilkina (2010). The Price Effect of Eliminating Potential Competition: Evidence from an Airline Merger. *Journal of Industrial Economics 58*(4), 767–793.

Lagos, V. (2018). Effectiveness of Merger Remedies: Evidence from the Retail Gasoline Industry. *Journal of Industrial Economics 66*(4), 942–979.

Li, L. and P. E. Strahan (2020). Who Supplies PPP Loans (And Does It Matter)? Banks, Relationships and the COVID Crisis. *Working Paper*.

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

Loertscher, S. and L. M. Marx (2021). Coordinated Effects in Merger Review. *Journal of Law and Economics*, Forthcoming.

Lopez, J. and M. Spiegel (2021). Small Business Lending Under the PPP and PPPLF Programs. *Working Paper*.

Luo, D. (2014). The Price Effects of the Delta/Northwest Airline Merger. *Review of Industrial Organization 44*(1), 27–48.

Mahoney, N. and E. Weyl (2017). Imperfect Competition in Selection Markets. *The Review of Economics and Statistics 99*(4), 637–651.

Majerovitz, J. and A. Yu (2021). Consolidation on Aisle Five: Effects of Mergers in Consumer Packaged Goods. *Working Paper*.

Manove, M., A. Padilla, and M. Pagano (2001). Collateral versus Project Screening: A Model of Lazy Banks. *RAND Journal of Economics 32*(4), 726–744.

Miller, N. H. and M. C. Weinberg (2017). Understanding the Price Effects of the MillerCoors Joint Venture. *Econometrica 85*(6), 1763–1791.

Mills, K. and B. McCarthy (2014). The State of Small Business Lending: Credit Access During the Recovery and How Technology May Change the Game. *Harvard Business School Working Paper, No. 15-004*.

Nelson, S. (2022). Private Information and Price Regulation in the US Credit Card Market. *Working Paper*.

Nocke, V. and N. Schutz (2018). Multiproduct-Firm Oligopoly: An Aggregative Games Approach. *Econometrica 86*(2), 523–557.

Nocke, V. and M. D. Whinston (2021). Concentration Thresholds for Horizontal Mergers. *American Economic Review*, Forthcoming.

Office of the Comptroller of the Currency (2015, 07). What is the SBA 7(a) Loan Guaranty Program? OCC Report.

Panetta, F., F. Schivardi, and M. Shum (2009). Do Mergers Improve Information? Evidence from the Loan Market. *Journal of Money, Credit, and Banking 41*(4), 673–709.

Peristiani, S. (1997). Do Mergers Improve the X-Efficiency and Scale Efficiency of U.S. Banks? Evidence from the 1980s. *Journal of Money, Credit, and Banking 29*(3), 326–337.

Pesendorfer, M. (2003). Horizontal Mergers in the Paper Industry. *RAND Journal of Economics 34*(3), 495–515.

Peters, C. (2006). Evaluating the Performance of Merger Simulation: Evidence from the US Airline Industry. *Journal of Law and Economics 49*(2), 627–649.

Phillips, R. (2013). Optimizing Prices for Consumer Credit. *Journal of Revenue and Pricing Management 12*(4), 360–377.

Pomatto, L., P. Strack, and O. Tamuz (2020). The Cost of Information. *Working Paper*.

Prager, R. A. and T. H. Hannan (1998). Do Substantial Horizontal Mergers Generate Significant Price Effects? Evidence from the Banking Industry. *Journal of Industrial Economics 46*(4), 433–452.

Rajan, U., A. Seru, and V. Vig (2015). The Failure of Models that Predict Failure: Distance, Incentives, and Defaults. *Journal of Financial Economics 115*(2), 237–260.

Ruckes, M. (2004). Bank Competition and Credit Standards. *The Review of Financial Studies 17*(4), 1073–1102.

Schmalensee, R. (2012). "On a Level with Dentists?" Reflections on the Evolution of Industrial Organization. *Review of Industrial Organization 41*(3), 157–179.

Scott Morton, F. (2019). Modern U.S. Antitrust Theory and Evidence Amid Rising Concerns of Market Power and its Effects: An Overview of Recent Academic Literature. Technical report, Center for Equitable Growth.

Shapiro, B. T., G. J. Hitsch, and A. E. Tuchman (2021). TV Advertising Effectiveness and Profitability: Generalizable Results from 288 Brands. *Econometrica 89*(4), 1855–1879.

Shapiro, C. (2021). Antitrust: What Went Wrong and How to Fix It. *Antitrust 35*(3), 33–45.

Simons, J. J. and M. Delrahim (2019). Hart-Scott-Rodino Annual Report: Fiscal Year 2019. Technical report, Bureau of Competition, Federal Trade Commission and Antitrust Division, Department of Justice.

Simpson, J. and C. Taylor (2008). Do Gasoline Mergers Affect Consumer Prices? The Marathon Ashland petroleum and Ultramar Diamond Shamrock Transaction. *The Journal of Law and Economics 51*(1), 135–152.

Starc, A. (2014). Insurer Pricing and Consumer Welfare: Evidence from Medigap. *RAND Journal of Economics 45*(1), 198–220.

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

Sweeting, A. (2010). The Effects of Mergers on Product Positioning: Evidence from the Music Radio Industry. *RAND Journal of Economics 41*(2), 372–397.

Train, K. (2009). *Discrete Choice Methods with Simulation* (2 ed.). Cambridge University Press.

Vita, M. and F. D. Osinski (2018). John Kwoka's Mergers, Merger Control, and Remedies: A Critical Review. *Antitrust Law Journal 82*(1), 361–88.

Wang, J. (2020). Screening Soft Information: Evidence from Loan Officers. *RAND Journal of Economics 51*(4), 1287–1322.

Wang, J. and D. H. Zhang (2021). The Cost of Banking Deserts : Racial Disparities in Access to PPP Lenders and their Implications. *Working Paper*.

Weinberg, M. (2008). The Price Effects of Horizontal Mergers. *Journal of Competition Law and Economics 4*(2), 433–447.

Weinberg, M. C. and D. Hosken (2013). Evidence on the Accuracy of Merger Simulations. *Review of Economics and Statistics 95*(5), 1584–1600.

Whinston, M. D. (2007). Antitrust Policy Toward Horizontal Mergers. In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization*, Volume 3, pp. 2369–2440. Elsevier.

Williamson, O. E. (1968). Economies as an Antitrust Defense: The Welfare Tradeoffs. *American Economic Review 58*(1), 18–36.

Yannelis, C. and A. Zhang (2021). Competition and Selection in Credit Markets. *Working Paper*.

Appendix A

Appendix to Chapter One

A.1 Event-Study Robustness

In this section, I display a number of event-study robustness results. Because the specifications shown in the main text rely on the selection of the relevant sample (i.e., the decision of which windows to examine), I first show that the results are robust to other window definitions. Table A.1 displays results for loan characteristics regressions when restricting to only loans issued within 14 days of a guarantee-rate change. The results from the main text are robust to this alternate definition.

Next, I examine whether the results in the main text can be explained by variation in the composition of lenders. Table A.2 displays results of the event-study regressions with lender fixed effects. These results provide the same qualitative conclusions as those in the main text, which suggests that the response to guarantee-rate increases does not act only through a change in lender participation. Conditional on participation in both the baseline and SBA Recovery period, when a lender receives a higher guarantee rate it offers loans with characteristics different from those in the baseline.

In Table A.3, I examine event-study results, restricted to loans issued around the two

	(1) Interest Rate (Pct.)	(2) Loan Return (Yearly)	(3) Amt. Borrowed ($ Thousands)	(4) Loan Size > 150,000	(5) Loan Term (Months)
Loans Issued Within 14 Days of Events					
SBA Recovery	-0.0706***	-0.0573***	60.24***	0.0934***	7.079***
	(0.0257)	(0.0115)	(19.12)	(0.0162)	(1.682)
Mean Outcome	5.87	2.12	586.87	0.83	168.42
Observations	5,027	5,027	5,027	5,027	5,027
Zip Code Dem. Controls	✓	✓	✓	✓	✓
Business Type FE	✓	✓	✓	✓	✓
NAICS (Two-Digit) FE	✓	✓	✓	✓	✓
Real Estate FE	✓	✓	✓	✓	✓
Event Date FE	✓	✓	✓	✓	✓

Standard errors are clustered by lender.
* p<0.1, ** p<0.05, *** p<0.01

Table A.1: Robustness to Window Size (Loan Characteristics)

Note: This table presents results of the event-study specifications, restricting to loans issued within 14 days of guarantee-rate changes. All specifications include controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date.

events in which the guarantee expansions lapsed. The American Recovery and Reinvestment Act of 2009 and the Small Business Jobs Act of 2010 included guarantee expansions as only one piece of a larger legislative agenda. To ease concerns that results are influenced by confounding variation from other aspects of the legislation, I show that the conclusions of the event-study analysis are robust to including only loans issued around the lapsation events. Because the expiration of the guarantee expansions occurred when funding ran out, it is unlikely that another, related policy change occurred contemporaneously.

Finally, I examine the robustness of results to the inclusion of price outliers. Table A.4 displays results for the main event-study specifications including observations for which prices lie below the first percentile or above the 99th percentile. Qualitatively, the takeaways are identical.

	(1) Interest Rate (Pct.)	(2) Loan Return (Yearly)	(3) Amt. Borrowed ($ Thousands)	(4) Loan Size > 150,000	(5) Loan Term (Months)
Loans Issued Within 42 Days of Events					
SBA Recovery	-0.0228	-0.0199***	41.39***	0.0545***	2.358**
	(0.0168)	(0.00629)	(10.16)	(0.00872)	(1.007)
Mean Outcome	5.85	2.12	558.66	0.81	165.41
Observations	13,466	13,466	13,466	13,466	13,466
Zip Code Dem. Controls	✓	✓	✓	✓	✓
Business Type FE	✓	✓	✓	✓	✓
NAICS (Two-Digit) FE	✓	✓	✓	✓	✓
Real Estate FE	✓	✓	✓	✓	✓
Event Date FE	✓	✓	✓	✓	✓
Lender FE	✓	✓	✓	✓	✓

Standard errors are clustered by lender.
* p<0.1, ** p<0.05, *** p<0.01

Table A.2: Robustness to Lender Fixed Effects (Loan Characteristics)

Note: This table presents results of the event-study specifications, including loans issued within 42 days of guarantee-rate changes. All specifications include lender fixed effects. Additionally, they include controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date.

	(1) Interest Rate (Pct.)	(2) Loan Return (Yearly)	(3) Amt. Borrowed ($ Thousands)	(4) Loan Size > 150,000	(5) Loan Term (Months)
Loans Issued Within 42 Days of Events					
SBA Recovery	-0.0417*	-0.0465***	81.15***	0.0922***	7.641***
	(0.0229)	(0.0102)	(14.61)	(0.0122)	(1.406)
Mean Outcome	5.85	2.12	575.51	0.83	166.37
Observations	7,277	7,277	7,277	7,277	7,277
Zip Code Dem. Controls	✓	✓	✓	✓	✓
Business Type FE	✓	✓	✓	✓	✓
NAICS (Two-Digit) FE	✓	✓	✓	✓	✓
Real Estate FE	✓	✓	✓	✓	✓
Event Date FE	✓	✓	✓	✓	✓

Standard errors are clustered by lender.
* p<0.1, ** p<0.05, *** p<0.01

Table A.3: Robustness to Including Lapse Events Only (Loan Characteristics)

Note: This table presents results of the event-study specifications, restricting to loans issued within 42 days of lapses of guarantee-rate expansions. All specifications include controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date.

	(1) Interest Rate (Pct.)	(2) Loan Return (Yearly)	(3) Amt. Borrowed ($ Thousands)	(4) Loan Size > 150,000	(5) Loan Term (Months)
Loans Issued Within 42 Days of Events					
SBA Recovery	-0.0582***	-0.0377***	57.02***	0.0566***	4.224***
	(0.0172)	(0.00753)	(10.64)	(0.00877)	(0.941)
Mean Outcome	5.85	2.13	557.55	0.80	164.02
Observations	14,278	14,278	14,278	14,278	14,278
Zip Code Dem. Controls	✓	✓	✓	✓	✓
Business Type FE	✓	✓	✓	✓	✓
NAICS (Two-Digit) FE	✓	✓	✓	✓	✓
Real Estate FE	✓	✓	✓	✓	✓
Event Date FE	✓	✓	✓	✓	✓

Standard errors are clustered by lender.

* $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table A.4: Event-Study Results – Price Outliers Included

Note: This table presents results of the event-study specifications, including loans issued within 42 days of guarantee-rate changes. All specifications include observations corresponding to price outliers, which were removed in the analyses in the main text. Additionally, the specifications include controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date.

A.2 Balance of Macroeconomic Indicators

The event-study framework relies on an assumption of common shocks across SBA guarantee periods. This appendix presents analyses to examine the balance of macroeconomic activity in the high- and low-guarantee periods. I estimate the following specification:

$$Y_{it} = \alpha_0 + \alpha_1 \mathbb{I}(t = SBA) + \epsilon_{it},$$

and I display results with and without event fixed effects. Standard errors are clustered by week in all specifications to account for serial correlation. I restrict to dates within 42 days of guarantee-rate changes and, to be consistent with the event-study specifications, I remove data from one week following the second take-up of more generous guarantees.

The variables of interest fall into three broad categories. First, I examine the balance of variables that underly bank lending costs, namely the Federal Funds Rate[1] and the one-month LIBOR rate based on the U.S. dollar.[2] The latter variable also enters the SBA's Fixed Base Rate, which determines interest rate caps. Second, I consider variables that capture broader U.S. economic activity. Namely, I analyze balance of the market yield on U.S. Treasury Securities at 10-Year Constant Maturity[3] and the market yield of U.S. Treasury Securities at 3-Year Constant Maturity.[4] Third, I examine stock prices for the three largest U.S. banks at the time: Bank of America, JPMorgan Chase, and Citigroup.[5]

[1]Source: Effective Federal Funds Rate, Federal Reserve Bank of New York, available at https://www.newyorkfed.org/markets/reference-rates/effr.

[2]Source: ICE Benchmark Administration Limited (IBA), 1-Month London Interbank Offered Rate (LIBOR), based on U.S. Dollar [USD1MTD156N], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/USD1MTD156N.

[3]Source: Board of Governors of the Federal Reserve System, Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity [DGS10], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/DGS10.

[4]Source: Board of Governors of the Federal Reserve System, Market Yield on U.S. Treasury Securities at 3-Year Constant Maturity [DGS3], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/DGS3.

[5]Source: Commodity Systems, Inc., retrieved from Yahoo Finance.

Results are displayed in Table A.5. We do not observe large differences in cost shifters and treasury yields, and most differences are not statistically significant. These results suggest that the adjustments to loan characteristics captured by the event studies are indicative of a response to the policy change rather than changes to lending costs. Bank stock prices are slightly higher in the SBA Recovery period, though the differences are marginally significant. These results are consistent with the policy change providing incentives for banking activity. In sum, the macroeconomic balance tests provide support for using the event-study framework to analyze the effect of guarantees on equilibrium loan characteristics and outcomes.

	No Event FE	Event FE	Mean
Federal Funds Rate	-0.006	-0.005	0.185
	(0.006)	(0.006)	
One-Month LIBOR	-0.009	-0.012*	0.336
	(0.027)	(0.007)	
Market Yield on U.S. Treasuries (10-Year Constant Maturity)	0.019	-0.001	3.048
	(0.104)	(0.052)	
Market Yield on U.S. Treasuries (3-Year Constant Maturity)	-0.020	-0.042	1.086
	(0.082)	(0.039)	
Bank of America Closing Stock Price	0.507	0.544	12.152
	(1.068)	(0.379)	
JPMorgan Chase Closing Stock Price	1.433	1.568*	36.624
	(1.996)	(0.818)	
Citigroup Closing Stock Price	2.139	2.269*	38.706
	(2.358)	(1.257)	

Standard errors are clustered by week.
* p<0.1, ** p<0.05, *** p<0.01

Table A.5: Macroeconomic Indicator Balance

Note: This table displays results of the macroeconomic balance specifications with (Column 1) and without (Column 2) event fixed effects. The time period of issue is the same as that for the main event-study specifications, namely within 42 days of guarantee-rate changes. Sources for each of the variables are noted in the main text.

A.3 Price and Ongoing-Fee Calculation

I model the loan price as the net present value of all loan payments, assuming full amortization. This measure captures the stated interest rate, as well as the loan repayment term. I assume a yearly-equivalent discount rate of 5%.

Denote r_{ij} the interest rate, T_{ij} the term (in months), and B_{ij} the amount borrowed. Under the assumption of full amortization, in any given month the remaining loan balance is

$$B_{ij} - (t-1)\frac{B_{ij}}{T_{ij}}$$

The monthly interest rate $\frac{r_{ij}}{12}$ is paid each period, meaning the aggregate payment over the course of the loan is given by

$$\sum_{t=0}^{T_{ij}} \frac{r_{ij}}{12}\left(B_{ij} - (t-1)\frac{B_{ij}}{T_{ij}}\right)$$

Assuming a monthly discount rate of $\delta = \frac{.05}{12}$ and normalizing by the size of the loan, the gross return is given by

$$R_{ij} = \frac{1}{B_{ij}} \frac{\sum_{t=0}^{T_{ij}} \frac{r_{ij}}{12}\left(B_{ij} - (t-1)\frac{B_{ij}}{T_{ij}}\right)}{(1+\delta)^t}$$

I then compute the yearly loan return as $y_{ij} = (1 + R_{ij})^{\frac{1}{T_{ij}/12}} - 1$. For the pricing regressions and structural analysis, I convert these yearly returns to a price for a normalized ten-year loan:

$$p_{ij} = (1 + y_{ij})^{10}$$

I compute the NPV of the ongoing guarantee fee using a similar procedure. The ongoing guarantee fee, f, per dollar guaranteed is paid each period. I again assume a monthly discount rate of $\delta = \frac{.05}{12}$ and compute the NPV of the ongoing fee as:

$$\psi(M_{ij}) = M_{ij}\frac{1}{B_{ij}}\frac{\sum_{t=0}^{T_{ij}} \frac{f}{12}\left(B_{ij} - (t-1)\frac{B_{ij}}{T_{ij}}\right)}{(1+\delta)^t}.$$

I convert to the equivalent for a normalized ten-year loan as above.

A.4 Preferred Lenders

As noted in the main text, preferred lenders are required to undergo screening by their SBA regional field office before being admitted to the program. Once approved, these lenders are given substantial independence in the lending process, and their loans are subject to less scrutiny than the loans of their counterparts. This streamlined application process suggests that preferred lenders are more able to adjust loan characteristics and respond to changes in policy, such as guarantee rates.

In this appendix, I examine heterogeneity in the event study and pricing analyses by preferred lender status. The findings suggest a differential response to the policy change by these pre-screened lenders. Namely, they decrease interest rates, increase loan amounts, and lengthen maturities by more than banks outside the preferred program. They also exhibit greater changes to the precision of risk pricing, indicating a potential difference in information acquisition cost across lenders.

I estimate the following regression for event-study specifications:

$$\begin{aligned}Y_{ijt} = \alpha + \delta_1 \mathbb{I}(t = \text{SBA Recovery}) + \delta_2 \mathbb{I}(j = \text{Preferred}) + \\ \delta_3 \mathbb{I}(t = \text{SBA Recovery}) \times \mathbb{I}(j = \text{Preferred}) + \beta X_{it} + \epsilon_{ijt}\end{aligned}$$

where, as in the main text, Y_{ijt} is the relevant loan characteristic (i.e., interest rate, yearly return, amount borrowed, loan term) for loan i issued by lender j in period t, and X_{it} is a vector of borrower covariates. These covariates include business type, NAICS (two-digit), and event-date fixed effects, as well as zip code demographics, which are the same covariates as those shown in the main text. Here, δ_3 is the coefficient of interest, measuring the incremental response of preferred lenders, relative to their counterparts. Results are displayed in Table A.6.

	(1) Interest Rate (Pct.)	(2) Loan Return (Yearly)	(3) Amt. Borrowed ($ Thousands)	(4) Loan Size > 150,000	(5) Loan Term (Months)
Loans Issued Within 42 Days of Events					
SBA Recovery	0.00118	-0.00788	26.19	0.0233*	0.716
	(0.0217)	(0.00971)	(16.71)	(0.0130)	(1.398)
Preferred Lender	-0.0501	-0.0264**	-114.8***	-0.0816***	3.231
	(0.0386)	(0.0124)	(18.36)	(0.0200)	(2.538)
SBA Recovery × Preferred Lender	-0.0516**	-0.0304***	60.49***	0.0602***	4.355**
	(0.0222)	(0.0111)	(20.39)	(0.0183)	(1.920)
Mean Outcome	5.85	2.13	558.23	0.81	164.33
Observations	13,992	13,992	13,992	13,992	13,992
Zip-Code Dem. Controls	✓	✓	✓	✓	✓
Business Type FE	✓	✓	✓	✓	✓
NAICS (Two-Digit) FE	✓	✓	✓	✓	✓
Real Estate FE	✓	✓	✓	✓	✓
Event Date FE	✓	✓	✓	✓	✓

Standard errors are clustered by lender.
* $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table A.6: Event Study Heterogeneity – Preferred Lender Status

Note: This table presents results of the heterogeneity analysis for the main event-study specifications, including loans issued within 42 days of guarantee-rate changes. All specifications include controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date.

For the pricing regressions, I estimate a two-stage framework similar to that in the main text separately for preferred and non-preferred lenders. Bootstrap distributions of interaction coefficients are displayed in Figure A.1.

The results indicate a clear differential response by preferred lenders. I now examine whether the lenders that participate in the program differ from those that do not. First, I focus on their geographic location. Figure A.2 displays the share of preferred lenders by state. These lenders are spread across geographies and appear not to cater to only a small number of markets.

While preferred lenders do not specialize in select geographic markets, they do tend to be larger than their counterparts. Large banks (i.e., those in the top quartile of the distribution of the quarterly average of total loans) are more likely to be preferred lenders than their

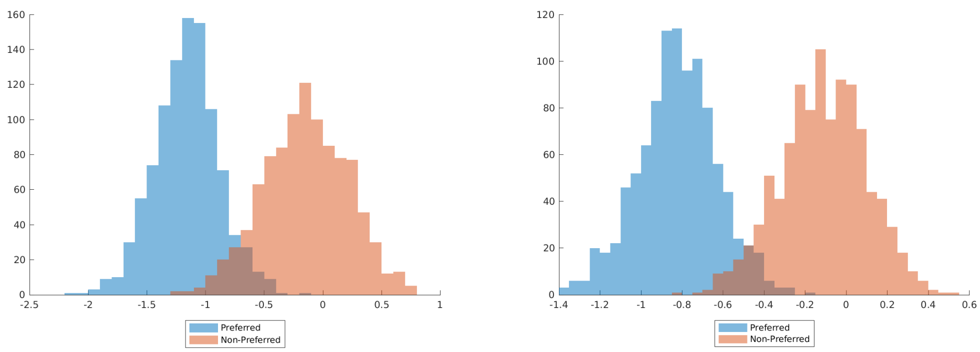

(a) Default (b) Charge Off

Figure A.1: Pricing Regression Heterogeneity – Preferred Lender Status

Note: This figure presents the bootstrap distribution of the coefficient on the interaction of the indicator for the SBA Recovery period and the pricing residual. This coefficient is estimated from the second stage of the pricing regression and is recovered separately for preferred and non-preferred lenders. It includes loans issued within 42 days of guarantee-rate changes and includes the same controls as used in the main text.

smaller peers. In all, 35.1% of banks in the top quartile of total loan value are preferred lenders, while this share is lower (17.0, 14.4, and 7.5%, respectively) for banks in the third, second, and first quartiles. Because of this, I show that the qualitative takeaways of the event-study results are the same when I restrict the sample to include only loans issued by banks in the top quartile of the distribution of total loans. Table A.7 displays estimates for the event-study specification with preferred-lender heterogeneity, restricted to loans issued by large banks. Furthermore, as mentioned in the main text, these lenders could differ along other, unobservable dimensions, so I allow information acquisition cost to flexibly vary across these categories of lenders.

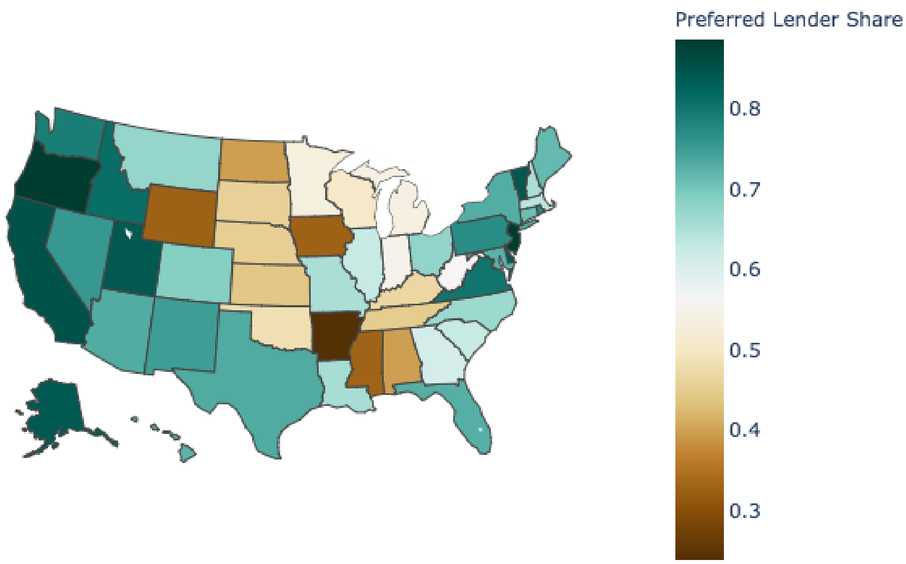

Figure A.2: Preferred Lender Share by State

Note: This figure displays a heat map of the share of preferred lenders by state. These shares are computed using all SBA 7(a) loans issued between 2009 and 2011.

	(1) Interest Rate (Pct.)	(2) Loan Return (Yearly)	(3) Amt. Borrowed ($ Thousands)	(4) Loan Size > 150,000	(5) Loan Term (Months)
Loans Issued Within 42 Days of Events					
SBA Recovery	0.0335	0.00633	31.97	0.0245	-0.217
	(0.0383)	(0.0154)	(24.75)	(0.0204)	(2.080)
Preferred Lender	-0.000749	0.00153	-150.4***	-0.109***	0.412
	(0.0549)	(0.0170)	(23.33)	(0.0264)	(3.425)
SBA Recovery × Preferred Lender	-0.0876**	-0.0445***	45.99*	0.0603**	4.678*
	(0.0345)	(0.0160)	(25.87)	(0.0262)	(2.693)
Mean Outcome	5.86	2.13	541.83	0.80	164.44
Observations	8,332	8,332	8,332	8,332	8,332
Zip Code Dem. Controls	✓	✓	✓	✓	✓
Business Type FE	✓	✓	✓	✓	✓
NAICS (Two-Digit) FE	✓	✓	✓	✓	✓
Real Estate FE	✓	✓	✓	✓	✓
Event Date FE	✓	✓	✓	✓	✓

Standard errors are clustered by lender.

* p<0.1, ** p<0.05, *** p<0.01

Table A.7: Event-Study Heterogeneity – Preferred Lender Status, Large Banks Only

Note: This table presents results of the heterogeneity analysis for the main event-study specifications, including loans issued within 42 days of guarantee-rate changes but restricting to lenders in the top quartile of the distribution of quarterly average of total loans. All specifications include controls for normalized zip code level demographics (median household income, total population, change in the housing price index since 2008), as well as fixed effects for business type, NAICS (two-digit), real estate, and event date.

A.5 Lender Competition

In this section, I provide evidence to support the assumption of lender market power in the structural model. The argument relies on one main institutional detail: the existence of the credit elsewhere test. For a borrower to receive funding through the SBA 7(a) program, they must not have outside funding options available. This detail implies that SBA lenders face limited, if any, competition from outside the program. For this reason, I focus on within-program competition in this section. Figure A.3 shows the distribution of unique lenders by borrower and the distribution of loan applications by borrower. The vast majority of borrowers are associated with only a single lender and apply for only a single guarantee between 2009 and 2011. While these facts do not rule out the existence of soft offers (i.e., borrowers receiving loan offers without first being approved for a guarantee), they do support the notion of limited competition.

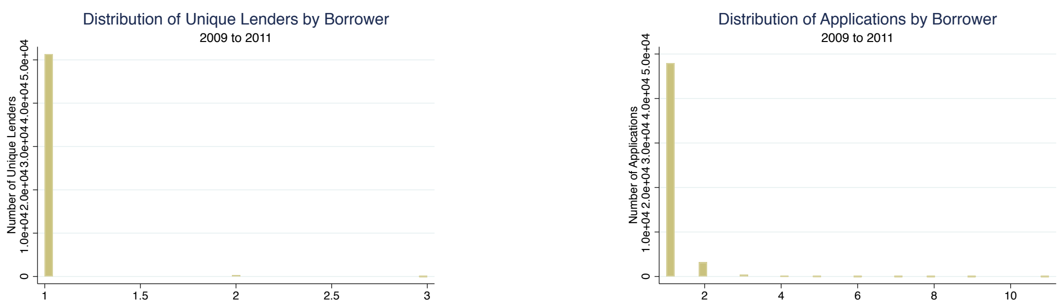

Figure A.3: Borrower Applications and Lender Relationships

Note: This figure displays the distribution of unique lenders by borrower (left-hand panel) and applications by borrower (right-hand panel) for all loans issued between 2009 and 2011.

To address the potential for soft offers, I now examine the relationship between prices and the number of bank branches in a given area. I estimate the following specification:

$$\log(p_{ij}) = \alpha \log(branch_{ij}) + \beta X_{ij} + \epsilon_{ij}$$

where $branch_{ij}$ is either the number of branches in a zip code, the number of unique branches in a zip code, or the same quantities in a county. X_{ij} are the same controls as used in the event studies. Results are displayed in Table A.8, and the relationship between prices and the number of branches is statistically significant, but not economically meaningful. This suggests that the effects of competition are limited.

	(1) $\log(p_{ij})$	(2) $\log(p_{ij})$	(3) $\log(p_{ij})$	(4) $\log(p_{ij})$
log(Branches Zip)	0.000872***			
	(0.000315)			
log(Unique Branches Zip)		0.000594*		
		(0.000340)		
log(Branches County)			-0.00177***	
			(0.000378)	
log(Unique Branches County)				-0.00266***
				(0.000602)
Observations	13,332	13,332	13,905	13,905

Standard errors are clustered by lender.

* p<0.1, ** p<0.05, *** p<0.01

Table A.8: Competition Regressions

Note: This table displays the results of the regressions examining the relationship between prices and measures of market concentration (i.e., numbers of branches and unique branches by geography). The loans included and controls are the same as those in the main event studies.

A.6 Maximum Likelihood Procedure

As discussed in the main text, I observe acceptance, default, and pricing decisions and seek to recover the vector Θ of model primitives. The likelihood function takes the following form:

$$\mathcal{L}\left((a_{ij}, d_{ij}, p_{ij}); \Theta\right) = \int\int\int \mathcal{L}\left((a_{ij}, d_{ij}, p_{ij})|\xi^R, \alpha, s; \Theta\right)\ dF_{\xi^R|\alpha,s}(\Theta)\, dF_{\alpha|s}(\Theta)\, dF_s(\Theta)$$

I now describe in detail the components of the likelihood and the procedure I use to find its maximum.

First, I decompose the joint likelihood into three components, corresponding to the separate default, acceptance, and pricing decisions.

$$\begin{aligned}
&\int\int\int \mathcal{L}\left((a_{ij}, d_{ij}, p_{ij})|\xi^R, \alpha, s; \Theta\right)\ dF_{\xi^R|\alpha,s}(\Theta)\, dF_{\alpha|s}(\Theta)\, dF_s(\Theta) \\
&\qquad = \int\int\int \mathcal{L}\left(d_{ij}|a_{ij}, p_{ij}, \xi^R, \alpha, s; \Theta\right) \mathcal{L}\left(a_{ij}|p_{ij}, \xi^R, \alpha, s; \Theta\right) \\
&\qquad\qquad \mathcal{L}\left(p_{ij}|\xi^R, \alpha, s; \Theta\right)\ dF_{\xi^R|\alpha,s}(\Theta)\, dF_{\alpha|s}(\Theta)\, dF_s(\Theta) \\
&= \int\int\int \mathcal{L}\left(d_{ij}|a_{ij}, \xi^R; \Theta\right) \mathcal{L}\left(a_{ij}|p_{ij}, \alpha; \Theta\right)\ dF_{\xi^R|\alpha,s}(\Theta)\, dF_{\alpha|s}(\Theta) \mathcal{L}\left(p_{ij}|s; \Theta\right)\ dF_s(\Theta)
\end{aligned}$$

Under the parametric assumptions described in the main text, I rewrite the likelihood of default and acceptance as follows:

$$\begin{aligned}
&\int\int \mathcal{L}\left(d_{ij}|a_{ij}, \xi^R; \Theta\right) \mathcal{L}\left(a_{ij}|p_{ij}, \alpha; \Theta\right)\ dF_{\xi^R|\alpha,s}(\Theta)\, dF_{\alpha|s}(\Theta) \\
&= \int\int \Phi\left(\frac{-X_i^R\beta^R - \mu_{\xi^R|\alpha,s}}{\sigma_{\xi^R|\alpha,s}}\right) \frac{\exp(X_i^A\beta^A - \alpha p_{ij})}{1 + \exp(X_i^A\beta^A - \alpha p_{ij})}\, dF_{\xi^R|\alpha,s}(\Theta)\, dF_{\alpha|s}(\Theta)
\end{aligned}$$

The likelihood of the price offer follows from the lender's first-order condition. Conditional on the signal, I back out the implied marginal cost error, ω_{ij}, which satisfies:

$$\omega_{ij} = \frac{\int \frac{\partial P^A}{\partial p}(\alpha, p_{ij}) \left(1-(1-M)P^D(\alpha)\right) p_{ij} + P^A(\alpha, p_{ij}) \left(1-(1-M)P^D(\alpha)\right) \, dF_{\alpha|s}}{\int \frac{\partial P^A}{\partial p}(\alpha, p_{ij}) \, dF_{\alpha|s}} - \beta^z z_{ij} - \zeta(M_{ij}),$$

where

$$P^A(\alpha, p_{ij}) = \frac{\exp(X_i^A\beta^A - \alpha p_{ij})}{1+\exp(X_i^A\beta^A - \alpha p_{ij})}, \; \frac{\partial P^A}{\partial p}(\alpha, p_{ij}) = -\alpha P^A(\alpha, p_{ij})(1-P^A(\alpha, p_{ij})), \text{ and}$$

$$P^D(\alpha) = \Phi\left(\frac{-X_i^R\beta^R - \mu_{\xi^R|\alpha,s}}{\sigma_{\xi^R|\alpha,s}}\right)$$

Then, I assemble the likelihood of the observed price offer as:

$$\mathcal{L}\left(p_{ij}|\xi^R, \alpha, s; \Theta\right) = \frac{1}{\sigma_\omega}\phi\left(\frac{\omega_{ij}}{\sigma_\omega}\right)\frac{\partial \omega_{ij}}{\partial p_{ij}}.$$

I obtain the maximum of the log-likelihood using the Interior/Direct Algorithm of Knitro and compute integrals using product-rule quadrature with 25 nodes in each dimension. I then obtain standard errors analytically, using the fact that the maximum likelihood estimator is asymptotically normal, where the distribution is given by:[6]

$$\sqrt{N}(\hat{\Theta} - \Theta) \to N(0, I(\Theta)^{-1}),$$

where $I(\Theta) = -E\left[\frac{\partial^2 \log \mathcal{L}(\Theta)}{\partial\Theta\partial\Theta'}\right]$. Under a correctly specified model, the information equality

[6] See, e.g., Greene (2012).

holds at the true value of the primitives, Θ. Thus,

$$I(\Theta) = -E\left[\frac{\partial^2 \log \mathcal{L}(\Theta)}{\partial\Theta\partial\Theta'}\right] = E\left[\frac{\partial \log \mathcal{L}(\Theta)}{\partial\Theta}\frac{\partial \log \mathcal{L}(\Theta)}{\partial\Theta'}\right],$$

and I estimate the standard errors using the variance-covariance matrix computed as the product of gradients of the log-likelihood function.

A.7 Full Results

In Table A.9, I present the full set of results for borrower- and lender-side primitives. The borrower-side estimates include the components of the utility of repayment, β^R, the components of the utililty of acceptance, β^A, and the primitives governing the joint distribution of borrower repayment and price responsiveness, μ_α, σ_α, and ρ. The lender-side primitives include the S.D. of the cost shock, σ_ω, the components of the marginal cost function, β^Z, and the S.D. of the signal distribution for different lender types and periods, $\sigma_\gamma(h_j)$.

Parameter	Estimate (S.E.)
Components of β^R:	
Constant	1.434***
	(0.215)
Normalized Amt. Borrowed	-0.084
	(0.095)
Maturity $\leq$ 10 Years	-0.419***
	(0.049)
Event 2	0.396***
	(0.061)
Event 3	0.302***
	(0.060)
Event 4	0.373***
	(0.055)
NAICS: 22	0.142
	(0.426)
NAICS: 23	-0.222
	(0.181)
NAICS: 31	-0.156
	(0.212)
NAICS: 32	-0.035
	(0.204)
NAICS: 33	-0.012
	(0.190)
NAICS: 42	0.091
	(0.190)
NAICS: 44	-0.196

	(0.175)
NAICS: 45	-0.264
	(0.190)
NAICS: 48	0.178
	(0.223)
NAICS: 51	0.231
	(0.271)
NAICS: 52	0.618*
	(0.347)
NAICS: 53	0.386
	(0.277)
NAICS: 54	0.076
	(0.182)
NAICS: 56	-0.181
	(0.196)
NAICS: 61	0.410
	(0.367)
NAICS: 62	0.149
	(0.177)
NAICS: 71	-0.371*
	(0.200)
NAICS: 72	-0.260
	(0.173)
NAICS: 81	-0.067
	(0.179)
Type: Individual	0.044
	(0.073)
Type: Partnership	0.209
	(0.130)
Normalized ΔHPI	0.230
	(0.225)
Normalized Median HH Inc.	0.285
	(0.176)
Normalized Tot. Pop	-0.066
	(0.124)
Components of β^A:	
Constant	22.310***
	(6.647)
Normalized Amt. Borrowed	-0.479***
	(0.168)

Maturity $\leq$ 10 Years	1.309***
	(0.455)
Event 2	-0.018
	(0.017)
Event 3	-0.021
	(0.017)
Event 4	-0.050**
	(0.023)
SBA Recovery	-0.041**
	(0.020)
Components of β^Z:	
Constant	0.654***
	(0.084)
Normalized Int. Bearing	0.153***
	(0.031)
Normalized Non-Int. Bearing	-0.144***
	(0.027)
Normalized Deposits	-0.043
	(0.028)
S.D. of Cost Shock, σ_ω:	
Constant	0.169***
	(0.034)
Components of $F_{\xi^R,\alpha}$	
μ_α	16.682***
	(4.987)
σ_α	0.701
	(1.032)
ρ	0.655
	(0.697)
S.D. of Signal Distribution: σ_γ	
Non-Preferred, Baseline	1.230***

	(0.108)
Preferred, Baseline	1.148***
	(0.109)
Non-Preferred, SBA Recovery	1.355***
	(0.104)
Preferred, SBA Recovery	1.485***
	(0.093)

Table A.9: Full Results – Lender-Side Parameters

A.8 Model: Random Coefficient on the Constant

In the main text, I present results for a model in which unobservable heterogeneity in demand takes the form of a random coefficient on price. Under this framework, adverse selection enters through the correlation of a borrower's unobservable utility of repayment and their price responsiveness. In this section, I show that results are robust to including the unobserved demand heterogeneity through a random coefficient on the constant term. In this case, adverse selection enters through the correlation of the unobservable utility of repayment and utility of acceptance, across all prices. The results presented in the main text are qualitatively the same as those obtained using this alternate specification.

As in the main text, the utility of loan repayment is given by:

$$u_i^R = X_i^R \beta^R + \xi_i^R,$$

where X_i^R is the same vector of borrower covariates as used in the main text and ξ_i^R is an unobservable shifter of the utility of repayment. Borrowers default if $u_i^R < 0$.

The expression for the utility of acceptance differs from the one in the main text and is instead:

$$u_{ij}^A = X_i^A \beta^A - \alpha p_{ij} + \xi_i^A + \epsilon_{ij},$$

where X_i^A is the same vector of acceptance shifters as in the main text and ξ_i^A is an unobservable shifter of the utility of acceptance. This unobservable is jointly distributed with the unobservable shifter of repayment, ξ_i^R. Specifically, I parameterize this joint distribution as:

$$\begin{pmatrix} \xi_i^R \\ \xi_i^A \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho\sigma_{\xi^A} \\ \rho\sigma_{\xi^A} & \sigma_{\xi^A}^2 \end{pmatrix} \right).$$

As in the main text, I make one assumption about the sign of ρ to ensure that prices are monotonic in signals. In this case, the non-advantageous selection assumption implies $\rho \leq 0$.

The lender's signal structure takes the same form as in the main text, and prices are set by solving:

$$\max_{p_{ij}} \int \int P^A(\xi_i^A, p_{ij}) \left[(1 - (1 - M) P^D(\xi_i^R)) p_{ij} - \zeta_{ij} \right] \, dF_{\xi_i^R, \xi_i^A | s_{ij}},$$

where the quantities are defined analogously to those in the body.

I now display updated versions of Tables 1.4 and 1.5 from the main text. The results for the model with a random coefficient on the constant are qualitatively the same as those with a random coefficient on price. Coefficients of the utility of repayment and utility of acceptance have the same sign and approximate magnitude as those in the main specification. Also, the joint distribution of risk and price responsiveness suggests a similar magnitude of heterogeneity and the presence of adverse selection, though the correlation parameter is again noisily estimated.

The results for the lender-side parameters are also similar. The standard deviation of the signal noise increases by 11.2% for preferred lenders (vs. 11.5% in the main text), while it decreases by 8.3% for non-preferred borrowers (vs. 8.4% in the main text). The combination of these results suggests that the main takeaways of the paper are robust to defining an acceptance decision with unobserved heterogeneity on the constant term rather than the price coefficient.

Parameter	Estimate (S.E.)
Components of β^R:	
Constant	1.415***
	(0.215)
Normalized Amt. Borrowed	-0.079
	(0.095)
Maturity $\leq$ 10 Years	-0.419***
	(0.049)
Event 2	0.400***
	(0.061)
Event 3	0.305***
	(0.059)
Event 4	0.377***
	(0.055)
NAICS: 52 (Finance & Insurance)	0.652*
	(0.352)
NAICS: 53 (Real Estate & Rental & Leasing)	0.390
	(0.276)
Components of β^A:	
Constant	22.234***
	(4.815)
Price	-16.664***
	(3.575)
Normalized Amt. Borrowed	-0.470***
	(0.109)
Maturity $\leq$ 10 Years	1.281***
	(0.291)
Event 2	-0.019
	(0.016)
Event 3	-0.022
	(0.016)
Event 4	-0.051***
	(0.019)
Components of F_{ξ^R,ξ^A}	
σ_{ξ^A}	0.719
	(1.129)
ρ	-0.739
	(0.995)

* p<0.1, ** p<0.05, *** p<0.01

Table A.10: Robustness – Selected Borrower-Side Estimates

Parameter	Estimate (S.E.)
S.D. of Signal Distribution: σ_γ	
Non-Preferred Lender, Baseline	1.163***
	(0.068)
Preferred Lender, Baseline	1.245***
	(0.069)
Non-Preferred Lender, SBA Recovery	1.066***
	(0.070)
Preferred Lender, SBA Recovery	1.384***
	(0.063)

* p<0.1, ** p<0.05, *** p<0.01

Table A.11: Robustness – Selected Lender-Side Estimates

A.9 Extensions: Hybrid Policy Counterfactuals

In the body of the paper, I present results that compare price offers and borrower surplus under (1) the baseline policy with a guarantee rate of 90% and no subsidy, (2) a hybrid policy with a guarantee rate of 50% and a subsidy set such that expected spending is the same as in policy (1), and (3) a policy with a guarantee rate of 50% and no subsidy. While comparing outcomes under these three policies illustrates the key forces at play when designing guarantee schemes in small-business lending markets, variation in rates and subsidies can be more flexible. In this section, I conduct two sets of analyses to show how prices and borrower surplus vary under alternative policies.

Figure A.4 considers hybrid policies with guarantee rates, $M \in \{0.5, 0.6, 0.7, 0.8\}$, and subsidies set such that expected spending is the same as that under a guarantee rate of 90% and no subsidy. As in the figures in the main text, I display results both in aggregate and broken down for borrowers in the top and bottom quintile of the repayment distribution. The takeaways from this analysis are qualitatively identical to those presented in the main text. The hybrid policy leads to gains in aggregate borrower surplus and tempers the distributional impact of the guarantee program.

The final extension considers alternative combinations of guarantees and subsidies, outside the fixed-spending pairs considered in Figure A.4. For this exercise, I display contour plots in Figure A.5 with guarantee rates on the horizontal axis and subsidies on the vertical axis. In the left-hand panel, lighter colors signify higher prices. In the righthand panel, the same lighter colors signify higher borrower surplus. The dashed line traces the iso-cost curve for the guarantee rate/subsidy pairs considered in the previous analysis.

These contour plots illustrate two key points. First, prices (borrower surplus) are decreasing (increasing) in both the subsidy level and the guarantee rate. Second, gains accrue more quickly when subsidies increase than when guarantees are more generous, which is con-

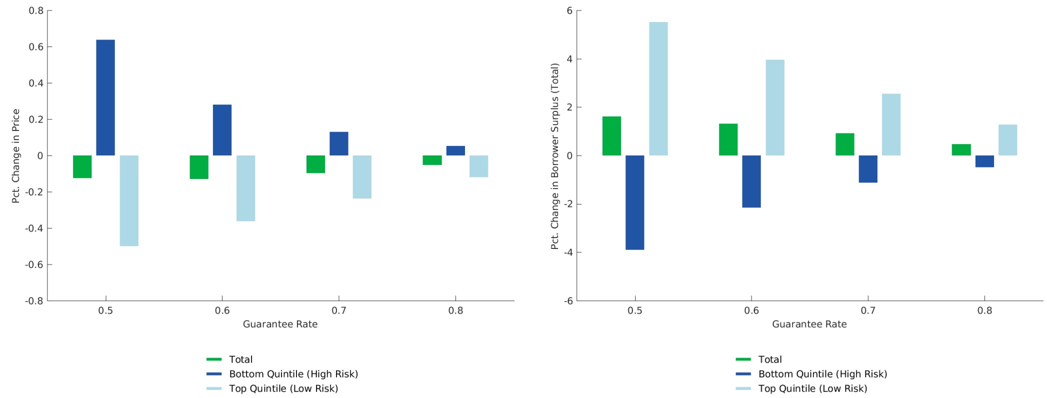

(a) Price (b) Borrower Surplus

Figure A.4: Changes to Price and Borrower Surplus Under Hybrid Policies

Note: This figure displays the average change in price and borrower surplus for (1) a hybrid policies with the specified guarantee rates and a subsidy set such that expected spending is the same as in the baseline (i.e., 90% guarantee and no subsidy).

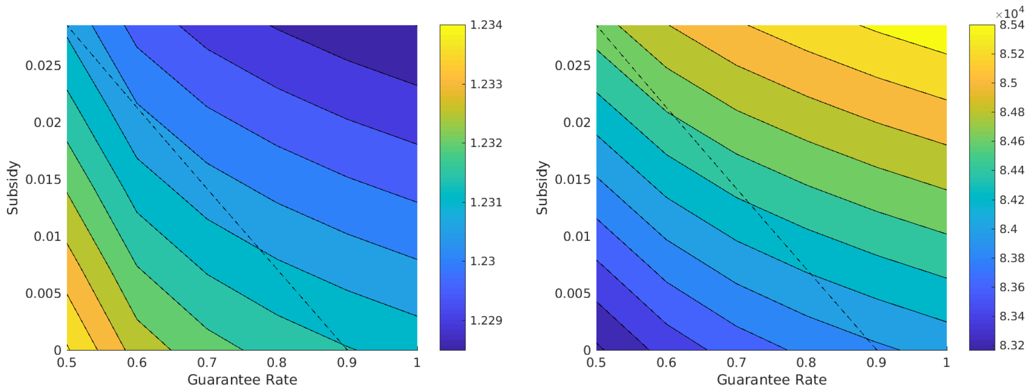

(a) Price (b) Borrower Surplus

Figure A.5: Price and Borrower Surplus Across Guarantees and Subsidies

Note: The contour plots display average prices and average borrower surplus, in $ over the normalized ten-year loan term, across guarantees and subsidies. The dashed line traces the iso-cost curve for the guarantee rates and subsidies considered in the analysis of the main text.

sistent with the main text's result that the hybrid policy yields higher aggregate borrower surplus.

Appendix B

Appendix to Chapter Two

B.1 Proofs

Theorem 1: Under Assumptions 1–2, for each θ_i, (i) there exists a unique $f^*(\theta_i) \in [0,1]$ that solves the borrower's optimization problem, and (ii) $f^*(\theta_i)$ is weakly decreasing in θ_i.

Proof: We begin by proving part (i), the uniqueness of f^*. Throughout, we condition on T, w, and b, and let $\underline{f}' = \frac{\underline{f}}{1+\frac{T\cdot w}{b}}$. There are five cases to consider:

1. $R(1,\theta_i) < 1$
2. $R(1,\theta_i) \geq 1$ and $1 + \frac{1-R(\underline{f}',\theta)}{\frac{\partial R}{\partial f}(\underline{f}',\theta)} > \frac{\underline{f}}{1+\frac{T\cdot w}{b}}$
3. $1 + \frac{1-R(\underline{f}',\theta)}{\frac{\partial R}{\partial f}} + \frac{1+r}{\gamma\cdot\underline{f}\cdot\frac{\partial R}{\partial f}} < \frac{\underline{f}}{1+\frac{T\cdot w}{b}}$
4. $1 + \frac{1-R(\underline{f}',\theta)}{\frac{\partial R}{\partial f}} + \frac{1+r}{\gamma\cdot\underline{f}\cdot\frac{\partial R}{\partial f}} \geq \frac{\underline{f}}{1+\frac{T\cdot w}{b}}$ and $1 + \frac{1-R(\underline{f}',\theta)}{\frac{\partial R}{\partial f}} \leq \frac{\underline{f}}{1+\frac{T\cdot w}{b}}$
5. $1 + \frac{1-R(0,\theta)}{\frac{\partial R}{\partial f}} + \frac{1+r}{\gamma\cdot\underline{f}\cdot\frac{\partial R}{\partial f}} < 0$

Consider case 1, where $R(1,\theta_i) < 1$. By Assumption 1, $R(\cdot,\theta_i)$ is increasing in its first argument. This implies that $R(f,\theta_i) < 1$ for all $f \in [0,1)$. Thus, $f^*(\theta_i) = 1$.

For the remaining cases, we must consider the borrower's first-order condition:

$$f_i^* = 1 + \frac{1 - R(f_i^*, \theta)}{\frac{\partial R}{\partial f}} + \frac{\mathbb{I}(f_i^*(1 + \frac{T \cdot w}{b}) < \underline{f})(1 + r)}{\gamma \cdot \underline{f} \cdot \frac{\partial R}{\partial f}}$$

For case 2, consider $f \in (\underline{f}', 1)$. First note that $1 + \frac{1-R(\underline{f}',\theta)}{\frac{\partial R}{\partial f}} > \frac{\underline{f}}{1+\frac{T \cdot w}{b}}$. We now must show that $G(f) = 1 + \frac{1-R(f,\theta)}{\frac{\partial R}{\partial f}}$ is strictly decreasing in f on $(\underline{f}', 1)$. We can show that

$$\frac{\partial G}{\partial f} < 0 \text{ if } (R(f, \theta) - 1)\frac{\partial^2 R}{\partial f^2} < \left(\frac{\partial R}{\partial f}\right)^2$$

which is satisfied if $R(f, \theta) - 1$ is log-concave. Because R is assumed to be concave in f, $R(f, \theta) - 1$ is also concave and, thus, log-concave in f. Therefore, by the intermediate value theorem, there exists a unique $f_i^* \in (\underline{f}', 1)$ that solves the borrower's problem. An analogous argument for case 3 can be used to prove a unique $f_i^* \in (0, \underline{f}')$.

Now, consider case 4. Monotonicity of $R(f, \theta)$ in its first argument and $1+\frac{1-R(\underline{f}',\theta)}{\frac{\partial R}{\partial f}}+\frac{1+r}{\gamma \cdot \underline{f} \cdot \frac{\partial R}{\partial f}} \geq \frac{\underline{f}}{1+\frac{T \cdot w}{b}}$ implies that, for all $f \in [0, \underline{f}')$, $1 + \frac{1-R(f,\theta)}{\frac{\partial R}{\partial f}} + \frac{1+r}{\gamma \cdot \underline{f} \cdot \frac{\partial R}{\partial f}} > f$. Furthermore, monotonicity and $1 + \frac{1-R(\underline{f}',\theta)}{\frac{\partial R}{\partial f}} \leq \underline{f}'$ implies that, for all $f \in (\underline{f}', 1]$, $1 + \frac{1-R(f,\theta)}{\frac{\partial R}{\partial f}} < f$. It follows that the borrower's problem is solved at $f^* = \frac{\underline{f}}{1+\frac{T \cdot w}{b}}$.

Finally, consider case 5. $1 + \frac{1-R(0,\theta)}{\frac{\partial R}{\partial f}} + \frac{1+r}{\gamma \cdot \underline{f} \cdot \frac{\partial R}{\partial f}} < 0$ and monotonicity of $R(\cdot, \theta)$ in its first argument implies that, for all $f \in (0, 1]$, $1 + \frac{1-R(f,\theta)}{\frac{\partial R}{\partial f}} < 0$. Therefore, the borrower's problem is solved at $f^* = 0$.

We conclude by proving part (ii) and show that $f^*(\theta)$ is weakly decreasing in θ. Denote the borrower's objective function as $u(f, \theta)$. It suffices to show that this objective function satisfies decreasing differences for values of θ such that an interior solution is optimal:

$$\frac{\partial^2 u}{\partial f \partial \theta} = b_i \gamma \left[-\frac{\partial R}{\partial \theta} + (1-f)\frac{\partial^2 R}{\partial f \partial \theta} \right] < 0,$$

where the inequality follows from $\frac{\partial R}{\partial \theta} < 0$ and $\frac{\partial^2 R}{\partial f \partial \theta} < 0$. ■

Theorem 2: Under Assumption 3, there exists a minimum loan amount, $\underline{b}$, such that lender j approves all loans with $b_i > \underline{b}$. The minimum loan amount is: (i) decreasing in S, (ii) decreasing in r, and (iii) increasing in $\underline{f}$. Specifically, $\underline{b}$ is given by:

$$\underline{b}(\tilde{f}_i^*) = \frac{c_{ij}}{S + (r - \delta_j)s(\tilde{f}_i^*)}$$

Proof: The lender earns the following profit from issuing a loan:

$$\pi_{ij} = Sb_i - c_{ij} + s(\tilde{f}_i^*)(r - \delta_j)b_i,$$

We normalize the lender's payoff of not issuing a loan to zero, and a loan is issued if $\pi_{ij} \geq 0$, which is satisfied if

$$b_i \geq \frac{c_{ij}}{S + s(\tilde{f}_i^*)(r - \delta_j)},$$

under the assumption that $S + s(\tilde{f}_i^*)(r - \delta_j) > 0$.

We now prove the comparative statics results:

$$\frac{\partial \underline{b}}{\partial S} = \frac{-c_{ij}}{(S + s(\tilde{f}_i^*)(r - \delta_j))^2} < 0$$

$$\frac{\partial \underline{b}}{\partial c} = \frac{1}{S + s(\tilde{f}_i^*)(r - \delta_j)} > 0$$

Consider a decline in the forgiveness threshold. For $\underline{f}'' < \underline{f}'$,

$$\begin{aligned}
\underline{b}(\underline{f}'') - \underline{b}(\underline{f}') &= \frac{c}{S + (r-\delta)s(\tilde{f}(\underline{f}''))} - \frac{c}{S + (r-\delta)s(\tilde{f}(\underline{f}'))} \\
&= \frac{c(r-\delta)(s(\tilde{f}(\underline{f}')) - s(\tilde{f}(\underline{f}'')))}{(S + (r-\delta)s(\tilde{f}(\underline{f}'')))(S + (r-\delta)s(\tilde{f}(\underline{f}'))} \\
&< 0
\end{aligned}$$

■

B.2 Window Definition

In this appendix, we conduct analysis to support the window definition used in the main text. For our main sample, we restrict to loans issued within the twelve week window encompassing the four weeks prior to the passage of the PPP Flexibility Act and the eight weeks following the policy change.

We use this sample for two main reasons. First, we seek to examine the lenders' response to the policy change and therefore seek to keep the distribution of borrowers approximately fixed. Our sample begins after the first month of the PPP. During this first month, discussions in the media and elsewhere highlighted the issues with the rollout of the program. Large businesses, such as Shake Shack, Potbelly Sandwich Shop, and Ruth's Chris Steakhouse, received loans[1], and borrowers faced numerous delays in receiving funds.[2] This difference in borrower types is apparent in the time series plot of loan amounts. Figure B.1 displays the 7-day moving average loan amount and total loan amount over time, and we see a large difference in the first four weeks of the PPP. To standardize the sample over time, we remove these loans issued at the beginning of the program.

Second, we restrict to loans issued up to eight weeks following the program because of a potential lag in policy take-up. In the plots below, we display the share of loans issued with an amount allocated to payroll of 75%, which was the threshold in the prior to the legislation change, and the share issued with an amount allocated to payroll of 60%, the threshold after the policy change. There is a transition period of about four weeks during which the share of loans issued at 75% decreases and the share issued at 60% increases. Because of this lag, we include an additional four weeks after which the policy take-up appears to stabilize. In Appendix B.3, we show the main empirical results are robust to

[1] https://www.nytimes.com/2020/04/20/business/shake-shack-returning-loan-ppp-coronavirus.html

[2] See, for example, https://www.usatoday.com/story/money/usaandmain/2020/04/07/ppp-loan-plan-rollout-disaster-small-businesses/2963901001/.

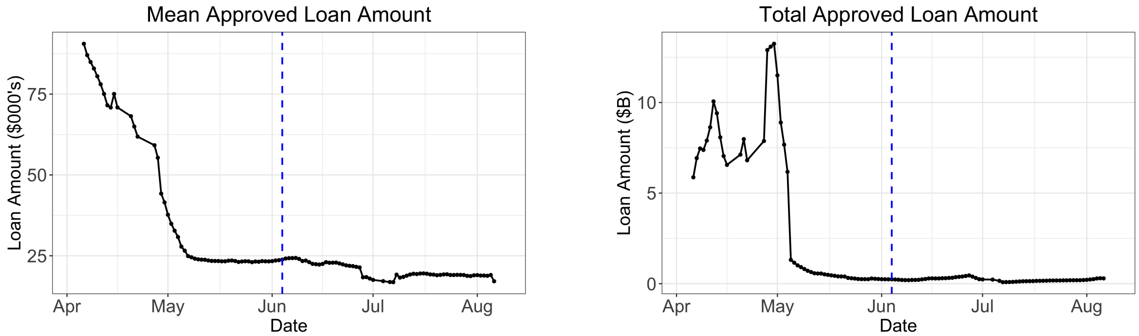

Figure B.1: 7-Day Moving Average of Mean and Total Loan Amount

Note: This figure displays two plots. On the lefthand side, we show the 7-day moving average of mean loan amount (in thousands of dollars) and, on the righthand side, we show the 7-day moving average of total loan amount (in billions of dollars).

using a donut-hole specification in which we remove the loans issued in the first four weeks following the policy change (i.e., the transition period).

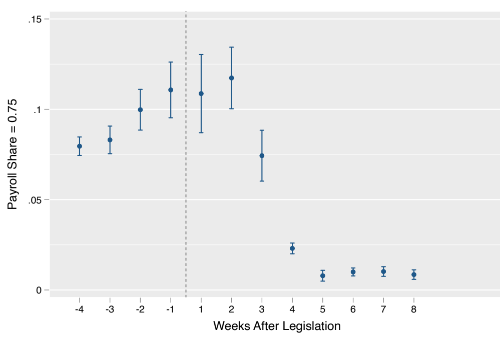

(a) Share = 0.75

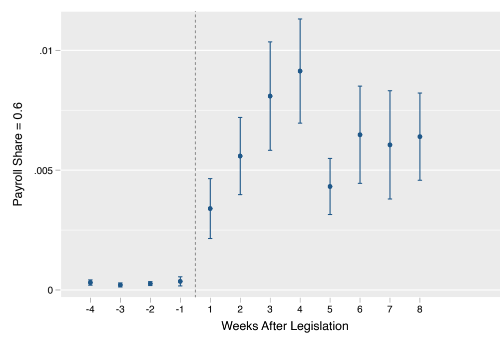

(b) Share = 0.60

Figure B.2: Share of Loans Issued at Threshold Payroll Amounts

Note: This figure the coefficient of a regression of an indicator of the at-origination payroll share being equal to 0.75 (lefthand side) or 0.6 (righthand size) on week fixed effects. The bars represent 95% conficence intervals computed using standard errors clustered by state.

B.3 Robustness Checks

Our main sample considers all loans issued between 4 weeks before and 8 weeks after the policy change. As described in Appendix , the take-up of the policy change occurred over time and took approximately four weeks to stabilize. For this reason, we examine robustness of our main event study results (i.e., in loan amount and in payroll share) to removing loans in the transition period. To do this, we restrict to loans issued up to four weeks before and between 5 and 8 weeks after the policy change.

Table B.1 displays the results of the event studies with loan amount on the lefthand side, and Table B.2 displays results for the payroll-share specifications. In both cases, the results are stronger for this specification than they are in the main text. However, the qualitative takeaways are identical.

	(1) Log(Loan Amount)	(2) Log(Loan Amount)	(3) Log(Loan Amount)
Post-Legislation	-0.128***	-0.139***	-0.0672***
	(0.0281)	(0.0217)	(0.0192)
Observations	661,062	660,808	660,805
Lender FE	No	Yes	Yes
Borrower Controls	No	No	Yes

* p<0.1, ** p<0.05, *** p<0.01

Table B.1: Aggregate Loan Amount Changes - Donut-Hole Specification

Note: This table presents results for the aggregate event-study specifications defined by Equation (2), restricting to loans issued up to four weeks before and between five and eight weeks after the policy change. Standard errors clustered by borrower state are shown in parentheses. Lender FEs are defined as a combination of a lender name, lender city, and lender state. Borrower controls include fixed effects for the two-digit NAICS code, business type (corporation, LLC, sole proprietorship, other), urban/rural, business age ($>$ 2 years, $\leq$ 2 years, unanswered), and borrower state.

	(1) Payroll Share (At Orig.)	(2) Payroll Share (At Orig.)
Post-Legislation	0.0184***	0.0174***
	(0.0020)	(0.0020)
Observations	661,062	661,055
Borrower Controls	No	Yes

* p<0.1, ** p<0.05, *** p<0.01

Table B.2: Aggregate Payroll Share Response - Donut-Hole Specification

Note: This table presents results for the event-study specifications with payroll share on the lefthand side, defined by Equation (3), restricting to loans issued up to four weeks before and between five and eight weeks after the policy change. Standard errors clustered by borrower state are shown in parentheses. Borrower controls include fixed effects for the two-digit NAICS code, business type (corporation, LLC, sole proprietorship, other), urban/rural, business age (> 2 years, ≤ 2 years, unanswered), and borrower state.

Appendix C

Appendix to Chapter Three

C.1 Additional Tables and Figures

In this appendix, we first provide additional tables containing information about the mergers we study. We document SIC codes of deals that were scrutinized and deals with overlap in Table C.1 and outline product market definitions in Table C.2.

We next provide some robustness analysis for the results presented in the body of the paper. Figure C.1 compares the estimated price and effects at the merger level across specifications. We show that the estimates are correlated at the merger level: not only is the distribution of price effects similar in the baseline to in the specifications with controls, but the estimates are similar at the merger level. This provides evidence that the subsequent analysis will not depend strongly on the fact that we use the baseline estimates. Table C.3 shows another dimension of robustness: we aggregate across mergers using the random effects weighting scheme proposed in DerSimonian and Laird (1986) instead of inverse variance weighting. Results are similar especially when taking into account standard errors, and the main difference is that the estimates for the specification comparing to untreated markets are more positive.

Figure C.2 replicates Figure 3.7 but using price changes for nonmerging parties. Here, the pattern is somewhat clearer that price changes are larger when ΔHHI is larger, while the correlation with HHI is not as large. Figure C.3 bins the scatters in Figures C.3 and C.2 to show average price effects by bin of HHI-ΔHHI.

Finally, Figure C.4 scatters the mean pre-merger time trend against the estimated price effects. We find no significant correlation, suggesting that price increases are not due to price trends.

SIC Code	Target Industry Sector	Deals Scrutinized	Deals with Overlap
212	Agriculture, Forestry, and Fishing	1	0
2011	Food and Kindred Products	1	1
2013	Food and Kindred Products	2	1
2015	Food and Kindred Products	2	1
2023	Food and Kindred Products	1	0
2026	Food and Kindred Products	1	0
2032	Food and Kindred Products	1	0
2033	Food and Kindred Products	3	2
2038	Food and Kindred Products	3	1
2041	Food and Kindred Products	1	0
2043	Food and Kindred Products	2	1
2047	Food and Kindred Products	1	1
2051	Food and Kindred Products	6	4
2052	Food and Kindred Products	1	0
2064	Food and Kindred Products	1	1
2068	Food and Kindred Products	1	1
2082	Food and Kindred Products	3	2
2084	Food and Kindred Products	4	1
2085	Food and Kindred Products	3	1
2086	Food and Kindred Products	3	1
2095	Food and Kindred Products	3	1
2096	Food and Kindred Products	3	2
2099	Food and Kindred Products	4	4
2111	Tobacco Products	1	0
2121	Tobacco Products	2	0
2621	Paper and Allied Products	1	1
2834	Drugs	4	0
2841	Soaps, Cosmetics, and Personal-Care Products	1	1
2844	Soaps, Cosmetics, and Personal-Care Products	9	6
2891	Chemicals and Allied Products	1	0
3841	Measuring, Medical, Photo Equipment; Clocks	1	0
5122	Wholesale Trade-Nondurable Goods	1	0
5143	Wholesale Trade-Nondurable Goods	1	1
5182	Wholesale Trade-Nondurable Goods	1	0
5431	Retail Trade-Food Stores	1	0
5461	Retail Trade-Food Stores	2	2
8049	Health Services	1	1

Table C.1: SIC Codes and Industry Sectors for Scrutinized Deals and for Deals in the Final Sample

Market	Nielsen Product Group	Nielsen Product Modules in Product Market
1	Baked Goods-Frozen	Bakery-Bagels-Frozen
2	Baked Goods-Frozen	Dough Products-Bread-Frozen, Bakery-Bagels-Frozen, Bakery - Doughnuts - Frozen, Bakery-Cheesecake-Frozen, Bakery - Biscuits/Rolls/Muffins - Frozen, Bakery-Breakfast Cakes & Sweet Rolls-Frozen, Bakery - Cobbler/-Dumplings/Strudel - Frozen, Bakery-Bread-Frozen, Bakery-Cookies Rte/-Cookie Dough-Frozen, Bakery - Dessert Cakes - Frozen, Bakery - Pies - Frozen, Bakery - Remaining - Frozen
3	Beer	Beer, Near Beer/Malt Beverage, Stout And Porter, Light Beer (Low Calorie/Alcohol), Ale, Malt Liquor
4	Beer	Beer, Stout And Porter, Light Beer (Low Calorie/Alcohol), Ale, Malt Liquor
5	Bread And Baked Goods	Bakery-Bread-Fresh
6	Bread And Baked Goods	Bakery-Bagels-Fresh
7	Bread And Baked Goods	Bakery-Breakfast Cakes/Sweet Rolls-Fresh
8	Bread And Baked Goods	Bakery-Buns-Fresh
9	Bread And Baked Goods	Bakery-Cheesecake-Fresh
10	Bread And Baked Goods	Bakery-Doughnuts-Fresh
11	Bread And Baked Goods	Bakery-Muffins-Fresh
12	Bread And Baked Goods	Bakery-Pies-Fresh
13	Bread And Baked Goods	Bakery-Rolls-Fresh
14	Candy	Candy-Chocolate-Miniatures, Candy-Chocolate, Candy-Chocolate-Special
15	Candy	Candy-Hard Rolled, Candy-Chocolate-Miniatures, Candy-Chocolate, Candy-Non-Chocolate-Miniatures, Candy-Non-Chocolate, Candy-Lollipops, Candy-Kits, Candy-Dietetic - Non-Chocolate, Candy-Dietetic - Chocolate, Gift Package With Candy Or Gum
16	Cereal	Cereal - Granola & Natural Types
17	Cereal	Cereal - Ready To Eat
18	Coffee	Coffee - Soluble Flavored, Coffee - Soluble
19	Coffee	Ground And Whole Bean Coffee, Coffee - Liquid
20	Condiments, Gravies, And Sauces	Barbecue Sauces
21	Condiments, Gravies, And Sauces	Cooking Sauce
22	Condiments, Gravies, And Sauces	Meat Sauce, Worcestershire Sauce
23	Condiments, Gravies, And Sauces	Sauce & Seasoning Mix-Remaining
24	Condiments, Gravies, And Sauces	Sauce & Seasoning Mix-Remaining Mexican
25	Condiments, Gravies, And Sauces	Sauce Mix - Spaghetti
26	Condiments, Gravies, And Sauces	Sauce Mix - Taco
27	Condiments, Gravies, And Sauces	Seasoning Mix - Chili
28	Condiments, Gravies, And Sauces	Seasoning Mix - Sloppy Joe
29	Cookies	Cookies
30	Cosmetics	Cosmetic Kits
31	Cosmetics	Cosmetics - Concealers
32	Cosmetics	Cosmetics-Blushers
33	Cosmetics	Cosmetics-Eye Shadows
34	Cosmetics	Cosmetics-Eyebrow & Eye Liner
35	Cosmetics	Cosmetics-Face Powder
36	Cosmetics	Cosmetics-Foundation-Liquid, Cosmetics-Foundation-Cream And Powder
37	Cosmetics	Cosmetics-Lipsticks
38	Cosmetics	Cosmetics-Mascara
39	Cosmetics	Cosmetics-Remaining
40	Crackers	Crackers - Flaked Soda, Crackers - Graham, Crackers - Sprayed Butter, Crackers - Cheese, Crackers - Remaining, Crackers - Flavored Snack, Snacks - Pork Rinds, Snacks - Puffed Cheese, Snacks - Potato Chips, Snacks - Potato Sticks, Snacks - Pretzel

Market	Nielsen Product Group	Nielsen Product Modules in Product Market
41	Detergents	Detergents-Packaged, Detergents - Light Duty, Detergents - Heavy Duty - Liquid
42	Detergents	Packaged Soap, Laundry Treatment Aids, Detergent Boosters, Fabric Washes - Special
43	Fragrances - Women	Cologne & Perfume-Women's
44	Grooming Aids	Cosmetic And Nail Grooming Accessory
45	Grooming Aids	Cosmetics - Noncotton Aplctrs/Puffs/Etc.
46	Gum	Gum-Bubble, Gum-Chewing, Gum-Chewing-Sugarfree, Gum-Bubble-Sugarfree, Breath Sweeteners
47	Hair Care	Creme Rinses & Conditioners
48	Hair Care	Hair Preparations - Other Than Men's
49	Hair Care	Hair Spray - Women's
50	Hair Care	Shampoo-Aerosol/ Liquid/ Lotion/ Powder, Shampoo-Combinations
51	Hair Care	Wave Setting Products
52	Jams, Jellies, Spreads	Garlic Spreads
53	Juice, Drinks - Canned, Bottled	Fruit Drinks-Canned, Fruit Drinks-Other Container, Water-Bottled
54	Liquor	Bourbon-Straight/Bonded, Bourbon-Blended, Canadian Whiskey, Irish Whiskey, Remaining Whiskey, Scotch, Gin, Vodka, Rum, Tequila, Brandy/Cognac, Cordials & Proprietary Liqueurs
55	Liquor	Vodka
56	Medications / Remedies / Health Aids	Foot Preparations-Athlete's Foot
57	Men's Toiletries	Cologne/Lotion-Men's
58	Packaged Meats-Deli	Bacon-Refrigerated, Sausage-Dinner, Sausage-Breakfast, Bacon-Beef & Canned
59	Packaged Meats-Deli	Bacon-Refrigerated
60	Packaged Meats-Deli	Bratwurst & Knockwurst, Frankfurters-Refrigerated, Franks-Cocktail-Refrigerated
61	Packaged Meats-Deli	Bratwurst & Knockwurst, Sausage-Dinner, Frankfurters-Refrigerated
62	Packaged Meats-Deli	Lunchmeat-Sliced-Refrigerated, Lunchmeat-Nonsliced-Refrigerated, Lunchmeat-Deli Pouches-Refrigerated
63	Packaged Meats-Deli	Lunchmeat-Sliced-Refrigerated
64	Packaged Meats-Deli	Sausage-Breakfast
65	Pet Food	Cat Food - Wet Type, Cat Food - Moist Type, Cat Food - Dry Type
66	Pet Food	Dog & Cat Treats
67	Pet Food	Dog Food - Wet Type, Dog Food - Moist Type, Dog Food - Dry Type
68	Pickles, Olives, And Relish	Pickles - Sweet
69	Pickles, Olives, And Relish	Relishes
70	Pizza/Snacks/Hors Doeurves-Frzn	Frozen/Refrigerated Hors D' Oeuvres & Snacks, Pizza-Frozen, Pizza Crust-Frozen, Meal Starters, Entrees - Remaining - 2 Food - Frozen, Entrees - Seafood - 2 Food - Frozen, Entrees - Meat - 2 Food - Frozen, Entrees - Poultry - 2 Food - Frozen, Entrees - Multi Pack - Frozen, Entrees - Italian - 2 Food - Frozen, Dinners-Frozen, Entrees - Mexican - 2 Food - Frozen
71	Prepared Food-Ready-To-Serve	Chicken - Shelf Stable
72	Prepared Food-Ready-To-Serve	Chili-Shelf Stable
73	Prepared Food-Ready-To-Serve	Stew - Beef - Shelf Stable, Stew - Remaining - Shelf Stable, Stew - Chicken - Shelf Stable
74	Prepared Foods-Frozen	Entrees - Meat - 1 Food - Frozen
75	Prepared Foods-Frozen	Entrees - Remaining - 2 Food - Frozen, Entrees - Seafood - 2 Food - Frozen, Entrees - Meat - 2 Food - Frozen, Entrees - Poultry - 2 Food - Frozen, Entrees - Multi Pack - Frozen, Entrees - Italian - 2 Food - Frozen, Dinners-Frozen, Entrees - Mexican - 2 Food - Frozen, Entrees - Meat - 1 Food - Frozen, Entrees - Poultry - 1 Food - Frozen, Entrees - Oriental - 1 Food - Frozen, Entrees - Italian - 1 Food - Frozen
76	Skin Care Preparations	Hand & Body Lotions
77	Skin Care Preparations	Hand Cream
78	Skin Care Preparations	Skin Cream-All Purpose
79	Snacks	Dip - Mixes
80	Snacks	Popcorn - Popped, Snacks - Caramel Corn

Market	Nielsen Product Group	Nielsen Product Modules in Product Market
81	Snacks	Snacks - Health Bars & Sticks
82	Snacks	Snacks - Potato Chips, Snacks - Potato Sticks
83	Snacks	Snacks - Potato Chips
84	Snacks	Snacks - Pretzel
85	Snacks	Snacks - Remaining
86	Spices, Seasoning, Extracts	Pepper
87	Spices, Seasoning, Extracts	Vegetables - Onions - Instant
88	Stationery, School Supplies	Dry Erase Bulletin Board And Accesory
89	Stationery, School Supplies	Personal Planners Binders And Folders
90	Unprep Meat / Poultry / Seafood-Frzn	Frozen Poultry
91	Vegetables - Canned	Mushrooms - Shelf Stable
92	Vegetables - Canned	Vegetables-Mixed-Canned
93	Vegetables - Canned	Vegetables - Peas - Remaining - Canned, Vegetables - Peas - Canned, Vegetables - Peas & Carrots - Canned

Table C.2: Product Market Definitions

	N	Mean	S.D.	25th Pct.	Median	75th Pct.
A. Baseline						
Overall	108	0.29	3.75	-2.15	-0.16	1.98
		(0.36)		(0.36)	(0.39)	(0.64)
Merging Parties	108	0.88	5.19	-2.07	1.03	3.97
		(0.49)		(0.74)	(0.51)	(0.74)
Non-Merging Parties	108	0.17	3.82	-2.13	-0.41	1.99
		(0.37)		(0.38)	(0.41)	(0.66)
B. Cost and Demographic Controls						
Overall	106	0.52	2.86	-0.92	0.01	1.74
		(0.27)		(0.40)	(0.27)	(0.52)
Merging Parties	106	1.22	4.84	-1.72	0.74	3.99
		(0.47)		(0.70)	(0.54)	(0.71)
Non-Merging Parties	106	0.46	2.85	-1.10	0.14	1.81
		(0.28)		(0.33)	(0.27)	(0.50)
C. Treated/Untreated						
Overall	101	0.64	8.18	-1.65	-0.01	1.21
		(0.78)		(0.36)	(0.24)	(0.42)
Merging Parties	101	1.40	9.33	-2.30	0.04	2.92
		(0.91)		(0.67)	(0.37)	(0.95)
Non-Merging Parties	101	0.52	7.91	-1.78	-0.28	1.09
		(0.74)		(0.34)	(0.27)	(0.40)

Table C.3: Overall Price Effects

Note: We aggregate across mergers using the random effects weighting scheme of DerSimonian and Laird (1986).

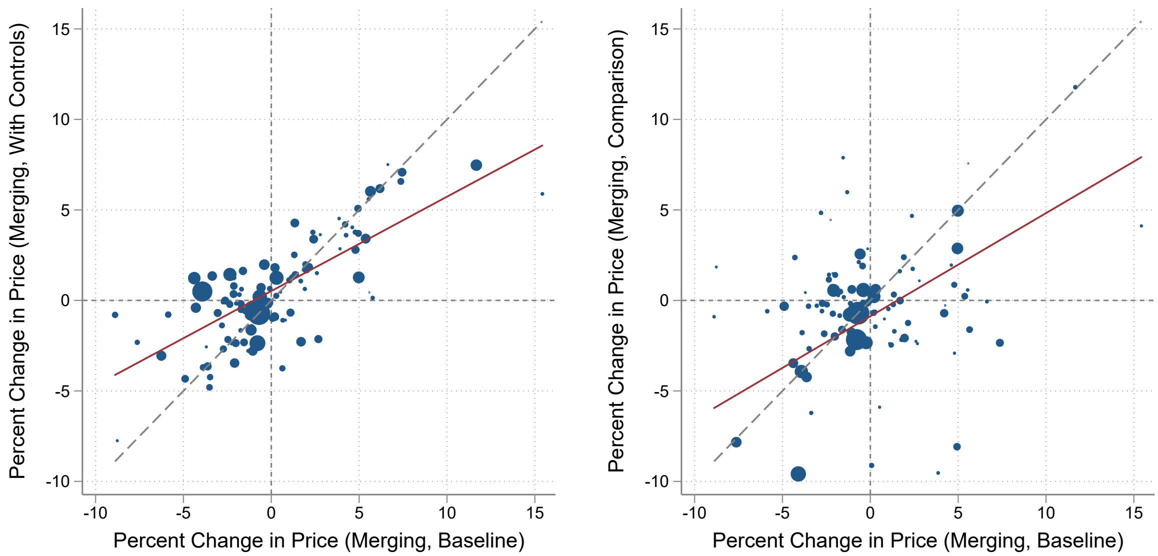

(a) Cost and Demographic Controls

(b) Untreated Markets

Figure C.1: Price Changes for Merging Parties, Estimated (a) with and without Controls and (b) Comparing to Untreated Markets or Not

Note: The size of each point corresponds to the inverse variance of the parameter estimate in the baseline.

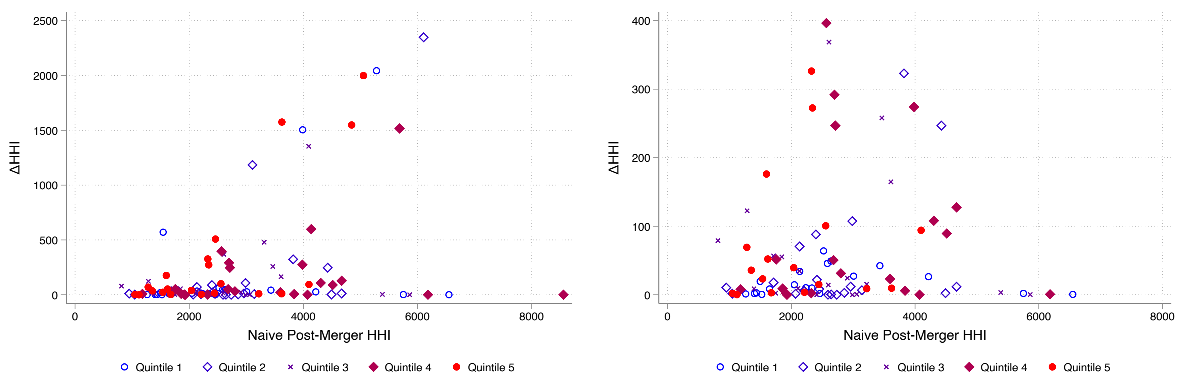

(a) Unrestricted

(b) Restricted

Figure C.2: Scatter Plots of Nationwide ΔHHI Against Nationwide Post-Merger HHI for Mergers in Our Sample

Note: Panel (a) displays the results for all mergers, while Panel (b) restricts to mergers for which $\Delta\text{HHI} < 400$ and post-merger $\text{HHI} < 7,000$. Colors/shapes represent the quintile of the price change for non-merging parties, with solid (red) colors representing larger price changes. Note that $\Delta\text{HHI} < 0$ corresponds to divestitures being required in the merger.

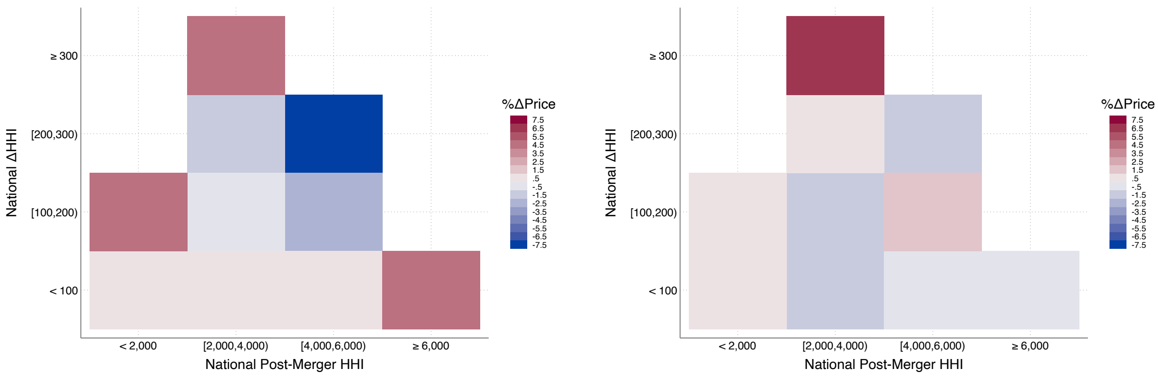

(a) Merging Parties (b) Non-Merging Parties

Figure C.3: Heat Map of Nationwide ΔHHI Against Nationwide Post-Merger HHI for Mergers in Our Dataset

Note: Colors represent different mean changes in price, as estimated by (3.1), for each bin of HHI and ΔHHI. Means are weighed by the inverse variance of each underlying parameter estimate.

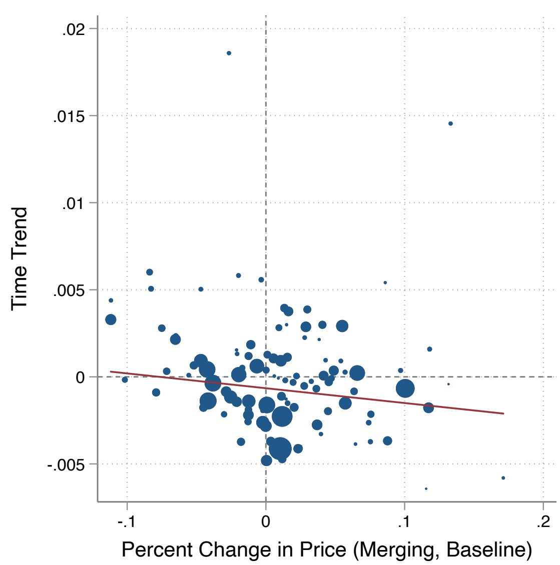

Figure C.4: Scatter Plot of Time Trend Estimated Using Only Pre-Merger Data versus Merging Party Price Effects as Measured by 3.1

Note: The size of each point corresponds to the inverse variance of the estimate of the price change.

C.2 Selection Model

In this appendix, we provide further details about the selection model in Section 3.6. Section C.2.1 discusses a general version of the model, and Section C.2.2 provides more detail of the estimation method.

C.2.1 Mechanics of the Model

Suppose that a merger is categorized by a vector (X_i, Z_i) of characteristics, and the distribution of price changes is given by $F(p_i, X_i)$. The antitrust authority proposes a remedy for a merger with probability $\lambda(p_i, X_i, Z_i)$, where Z_i has sufficient variation that for sufficiently extreme Z_i $\lambda(p_i, X_i, Z_i)$ is arbitrarily close to 0. Importantly, Z_i enters into the probability of a remedy but not into the distribution of price changes.

Suppose that true price changes are observed if the merger is allowed, and if the merger is not allowed we know that it is not. For a sufficiently extreme Z_i, we will observe the true

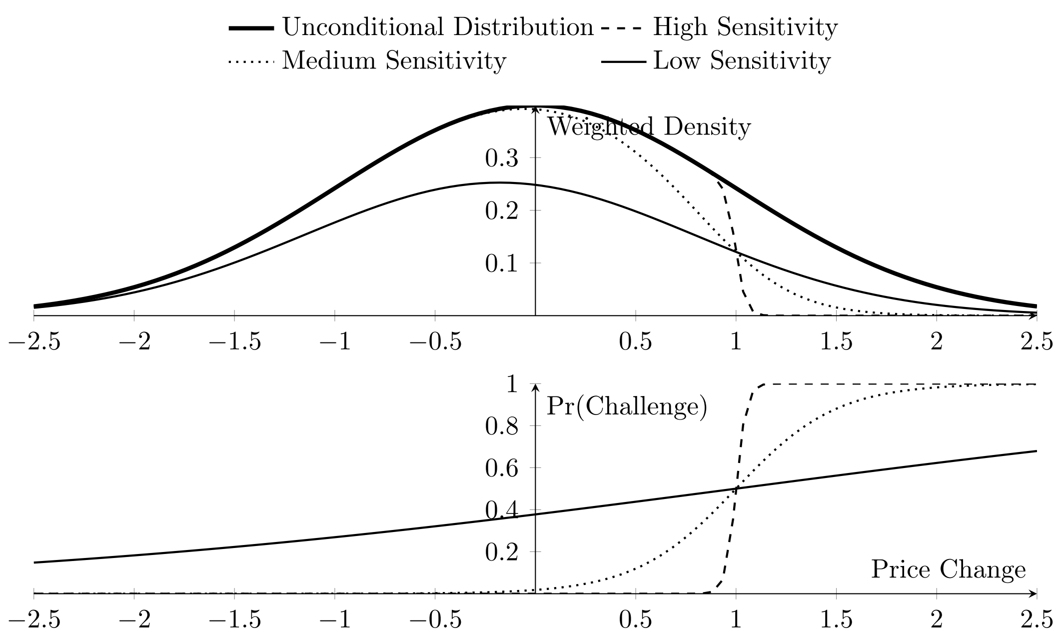

Figure C.5: Observed Distribution of Prices (top) and Implied $\lambda(\cdot)$ (bottom) distribution $F(X_i)$ of price changes. This is illustrated by the bold dark distribution in the top panel of Figure C.5. For different values of Z_i, we observe the density of price changes of consummated mergers too—as well as the probability a merger was not challenged. This give rise to a set of "weighted densities" illustrated in the top panel—distributions that integrate to the probability of consummation (without remedies) given Z_i. The thinner lines illustrate three possibilities of such weighted densities, depending on the agencies' enforcement behavior. All these densities lie beneath the unconditional density.

Everything in the top panel is data. The ratio of the unconditional density and the weighted one gives $\lambda(p_i, X_i, Z_i)$, illustrated for the three possibilities. The dashed line (labeled "high sensitivity") is a situation where the agency challenges mergers with price changes above 1% and lets others through. The "medium" and "low" sensitivity lines are policies where the agency challenges some mergers with lower price increases as well and is not as aggressive at challenging mergers with larger price changes.

The function $\lambda(p_i, X_i, Z_i)$ is an estimable primitive. In Section 3.6, we interpret the agency's decision as a threshold: it challenges a merger if its perceived price increase is larger than some cutoff $\bar{p}(X_i, Z_i)$. This interpretation is consistent with the model here. Since $\lambda(\cdot)$ is increasing in p_i and bounded in $[0, 1]$, it is a CDF of some random variable $\tilde{\epsilon}(X_i, Z_i)$. Thus, we can write

$$\begin{aligned}
\lambda(p_i, X_i, Z_i) &= \Pr(\tilde{\epsilon}(X_i, Z_i) \leq p_i) \\
&= \Pr(\tilde{\epsilon}(X_i, Z_i) - \mathbb{E}[\tilde{\epsilon}(X_i, Z_i)] \leq p_i - \mathbb{E}[\tilde{\epsilon}(X_i, Z_i)]) \\
&= \Pr(p_i - \epsilon_i \geq \mathbb{E}[\tilde{\epsilon}(X_i, Z_i)]), \qquad \text{(C.1)}
\end{aligned}$$

where ϵ_i is the demeaned version of $\epsilon(X_i, Z_i)$. The calculations in (C.1) show that $\bar{p}(X_i, Z_i)$ is simply the expectation of the variable whose cdf is given by the bottom panel of Figure C.5. When this function is close to an indicator for being large than a threshold (in true price changes), the associated expectation is of course this threshold. In general cases, it can be interpreted as an average of prices, weighted by how price-sensitive the probability of a challenge is at that true price.

The agencies likely have more information about mergers. Suppose the agencies make their decisions based on X_i, Z_i and some merger-specific information $\tilde{X}_i$ not available to us. In this case, what we recover is simply the expected probability of challenging a merger with a price change of p_i and observable characteristics X_i and Z_i, i.e.,

$$\lambda(p_i, X_i, Z_i) = \int \Pr(\text{challenge}|p_i, X_i, \tilde{X}_i, Z_i)\ dF(\tilde{X}_i|X_i).$$

The expected value of the random variable with this cdf will still be interpreted as an average threshold. However, this observation highlights that the error term should not be interpreted literally as an uncertainty in the estimate of a price change; it could be masking other considerations not taken into account into the analysis without necessarily threatening the interpretation of $\bar{p}$ as a threshold.

C.2.2 Estimation Details

The model discussed in Section 3.6 is a parametric version of the one in Appendix C.2.1, and it lends itself to a Bayesian estimation procedure involving Gibbs sampling with data augmentation; in particular, we augment with the true price change p_i^* (which is unobserved) and the belief errors ϵ_i. This involves the following steps.

1. *Initialize.* We first initialize variables that we do not observe. This involves picking $\beta_{\bar{p}} = 0$, setting $\epsilon_i = 0$, and then setting $p_i^* = -1$ for each merger that was approved without a remedy and $p^* = 1$ for all other mergers.

2. *Draw p_i^*.* For mergers that had a remedy, all we know is that $p_i^* \sim N(X_i'\beta_{p^*}, \sigma_{p^*})$,

truncated to being above $\bar{p}(X_i) - \epsilon_i$. For mergers that were allowed through, we know that $p_i^* \sim N(\tilde{\mu}, \tilde{\sigma}^2)$, where $\tilde{\mu}$ and $\tilde{\sigma}$ are the appropriate mean and standard deviation of the conditional distribution given the observed price change (with noise) is p_i.

3. *Draw β_{p^*} and σ_{p^*}.* This is Bayesian OLS of p_i^* on X_i.

4. *Update ϵ_i.* This involves drawing from truncated normals, knowing that $p_i^* + \epsilon_i$ must be below or above $\bar{p}(X_i)$, depending on whether the merger was let through without remedies.

5. *Draw $\beta_{\bar{p}}$.* This involves a Gibbs sampler through each dimension. The conditional distribution of each $\beta_{\bar{p},i}$ is truncated normal, where the truncation points are picked to restrict the constraints that $p_i^* + \epsilon_i \leq \bar{p}(X_i)$ for all mergers that were allowed without remedies, and the opposite inequality holds for all other mergers. We then loop back to Step 1.