

Indirect Effects of Access to Finance[†]

By JING CAI AND ADAM SZEIDL*

We created experimental variation across markets in China in the share of firms having access to a new loan product. Access to finance had a large positive direct effect on the performance of treated firms but a similar-sized negative indirect effect on that of firms with treated competitors, leading to nondetectable gains in producer surplus. Access to finance had a positive direct effect on business quality and consumer satisfaction and a negative effect on price, which were not offset by indirect effects, implying net gains in consumer surplus. We document other indirect effects and combine effects in a welfare evaluation. (JEL D22, G21, G32, L81, O16, P31, P34)

Lack of credit is widely believed to be a growth barrier, and a growing literature documents that credit has positive effects on firms' performance. However, in industry equilibrium, credit can have indirect effects on other actors: on peer firms through business stealing, information diffusion, and other channels, and on consumers through improvements in borrower firms' quality and price. Such industry equilibrium effects can be important for welfare and redistribution, but at present, we have limited evidence on their nature and magnitude. Rotemberg (2019) documents a negative indirect effect of a subsidy policy on peer firms in India, and McKenzie and Puerto (2021) document no indirect effect of a microenterprise training experiment on peer firms in Kenya. But we know little about the indirect effects of credit on both firms and consumers, the underlying mechanisms, and the welfare implications.

To investigate these issues, we conducted an experiment with 3,173 firms in 78 retail markets in China, in which we created cross-market variation in the share of firms receiving support in accessing a new loan product. This design allows us to identify both direct and indirect effects. We show that the loan had a large direct effect on the revenues and profits of treated firms. However, this was largely offset

*Cai: Department of Agricultural and Resource Economics, University of Maryland (email: cai516@umd.edu); Szeidl: Department of Economics, Central European University (email: szeidla@ceu.edu). Rema Hanna was the coeditor for this article. We thank David Atkin, Abhijit Banerjee, Lauren Bergquist, Nick Bloom, Arun Chandrasekhar, Shawn Cole, Dave Donaldson, Sean Higgins, Jonas Hjort, Dean Karlan, Miklos Koren, David McKenzie, Timea Molnar, Ben Olken, Martin Rotemberg, Antoinette Schoar, Eric Verhoogen, Chris Woodruff, and seminar and conference participants for helpful comments. We are grateful for funding from the Private Enterprise Development in Low-Income Countries, the University of Michigan, Central European University, the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) grant agreement number 283484, and the European Research Council under the European Union's Horizon 2020 research and innovation programme grant agreement number 724501. University of Michigan claims IRB approval is not required for the research in this paper (Study eResearch ID: HUM00087864). This study was registered in the American Economic Association Registry for randomized control trials under trial number AEARCTR-0009506.

[†]Go to <https://doi.org/10.1257/aer.20220711> to visit the article page for additional materials and author disclosure statement(s).

by a similar-sized business-stealing effect on competitors, implying that the program had no detectable effect on producer surplus. Turning to mechanisms, we show that treated firms improved quality and reduced price and that consumers valued these changes, implying that the program had a large positive effect on consumer surplus. We document additional indirect effects driven by the diffusion of information and the diffusion of demand. We then combine the direct and indirect effects in a welfare analysis and show that indirect effects are important for understanding industry outcomes and evaluating industrial policy.

In Section I, we present our context and experimental design. In 2013, a large bank introduced a new loan product to small and medium-sized enterprises (SMEs) in a province located in southeastern China. The loan product was offered in local “markets”—clusters of retail and service firms—and provided better terms than existing alternatives, primarily in that it did not require collateral. In collaboration with the bank, in the summer of 2013, we introduced a randomized treatment to improve access to the new loan product: we had a loan officer visit every treated firm once a month for a year, explain the terms of the loan, and offer help in filling out the application.

We randomized treatment intensity across 78 markets. In 37 markets we treated 80 percent of the firms, in 10 markets we treated 50 percent of the firms, and in 31 markets we treated no firms. This design generated random variation in both the firm’s treatment status and the share of its peers that were treated. We complemented the randomized intervention with surveys of our sample of 3,173 firms (average employment of 9). We conducted long surveys in 2013, before the intervention (baseline); in 2015, which was two years after the intervention (midline); and in 2016 (endline). We also conducted a short follow-up survey in 2020 to collect data on prices, consumer satisfaction, and other outcomes that were not included in our long surveys.

We begin our analysis in Section II with a simple model of the impacts of borrowing that incorporates two types of indirect effects: information diffusion and business stealing. In the model, firms compete monopolistically in local markets. To capture the idea that the loan enables productive investments, we assume that borrowing leads to higher product quality (or variety) and lower marginal cost. We allow information about the loan to diffuse to untreated firms, and model their decision to borrow. The main prediction of the model is that the impact of the loan program on (log) revenue can be approximated with a linear function that (i) depends positively on the borrowing status of the firm, representing a direct effect; and (ii) depends negatively on the share of the firm’s competitors that borrow, representing a business-stealing effect. Information diffusion about the loan affects performance only through borrowing and hence does not enter the equation. Because borrowing is endogenous, this equation cannot be directly applied to the data, but we show that it can be estimated with an instrumental variables (IV) strategy in which the treatment and the share of competitors treated are the instruments. The reduced form of this IV motivates our main empirical specification, in which the key explanatory variables are the firm’s treatment status and the share of the firm’s competitors that are treated.

In Section III, we present our reduced-form empirical results. We begin by looking at take-up. We find that, by endline, 28 percentage points more treated than untreated firms borrowed using the new loan product. The new loan did not crowd out

other borrowing, indicating that firms had genuine borrowing constraints (Banerjee and Duflo 2014). Most important, there was a positive indirect effect on borrowing: the share of peer firms in the market that were treated had a positive effect on the borrowing of untreated firms. This result is consistent with the information diffusion mechanism of our model.

We then look at direct and indirect effects on firm performance, using the empirical specification implied by the model. When the outcome is log sales, we estimate a significant direct effect of the treatment of about 10 log points and a significant indirect effect of the share of competitors treated of about -9 log points. The similar magnitude of these coefficients suggests a small impact on market-level revenue and indicates that the treatment induced a reallocation of revenue from firms with treated competitors to treated firms. Similarly, we find a large positive direct and a similar-sized negative indirect effect on profit, suggesting a small impact on market-level producer surplus. Turning to factors and inputs, we find a similar reallocation for employment and the wage bill, a smaller and less significant reallocation for materials, and—consistent with these firms having low capital intensity—no effect on fixed assets. Overall, the results on these main outcomes are consistent with the model's prediction on business stealing.

We then explore the mechanisms behind these results. In our model, business stealing is driven by two mechanisms (see Figure 1): (i) improvements in quality and variety and (ii) reductions in cost and price, both of which increase consumers' satisfaction with borrowers and induce a reallocation of demand to borrowers. We begin by documenting that the treatment had a positive direct and negative indirect effect on the number of clients, consistent with the demand reallocation interpretation. We then turn to the two mechanisms. We first consider firms' use of business practices that may improve quality and variety: renovation, the introduction of new products, and the share of workers with a high school degree. For all three, we estimate a large and significant direct effect and a small and less significant indirect effect. We then consider firms' use of practices that may reduce cost: switching to a new supplier, the stocking period (the average time between restocking, positively related to the amount of stock and hence the supplier's discount) and the use of inventory management techniques. For all three, we estimate a large and significant direct effect and an insignificant indirect effect. We conclude that the evidence is consistent with the model's two mechanisms.

We then investigate whether, as Figure 1 predicts, the changes in practices improved consumer satisfaction. To do this, we turn to measures of consumer experience. In our 2020 follow-up survey, we collected data on price and on various dimensions of consumer experience. We measured the latter with the evaluations of a randomly chosen customer, which is an approach that may be useful in other retail and service contexts as well. For price, we estimate a negative direct effect and no indirect effect. For each dimension of consumer experience, and for an index of overall consumer satisfaction, we estimate a positive direct effect and a much smaller and less significant indirect effect. We conclude that consumers valued the improved practices enabled by the loan, consistent with the mechanisms driving demand reallocation in the model.

We then turn to market-level effects, which can be estimated directly by regressing market-level outcomes on the share of firms treated in the market. We find



FIGURE 1. MECHANISMS DRIVING DEMAND REALLOCATION

insignificant effects on market revenue or profit but significant improvements in market-level business quality and consumer satisfaction. These results closely correspond to those obtained by aggregating the firm-level results. Our model can rationalize these patterns with a low market-level elasticity of demand (near the Cobb-Douglas level), which implies that despite improvements in price-adjusted quality the market-level expenditure remains essentially constant.¹

Our analysis thus far has assumed that indirect effects affect all peers with equal intensity. We next explore the heterogeneity of these effects by distance and competition status. First, we consider indirect effects on borrowing. We show that the diffusion of borrowing is coming from similar peers not directly competing with the firm: neighboring noncompetitors and distant competitors. There is no diffusion from neighboring competitors, suggesting that they withhold information to avoid business-stealing effects (Cai and Szeidl 2018; Hardy and McCasland 2021). Next, we consider indirect effects on firm performance. Our main result is that treating noncompetitor neighbors has a *positive* effect. We present various pieces of evidence that this effect is not explained by the diffusion of borrowing and conclude that it represents a novel indirect effect: the diffusion of demand from treated stores to noncompetitor neighbors, plausibly driven by customers shopping around in the neighborhood. This result provides experimental evidence on a demand externality that may drive the spatial concentration of retail observed around the world as well as in our markets (Leonardi and Moretti 2022).

In Section IV, we combine the direct and various indirect effects in a model-based welfare evaluation. We first estimate the IV regression implied by the model to infer the direct and business-stealing effects of the loan. Paralleling the reduced-form results, we find significant and largely offsetting effects on revenue and profit. We then use these estimates to quantify the impact of the loan program on welfare and redistribution. Our model implies that we can infer the gain in consumer surplus from the direct effect of the loan on firm revenue, which measures the extent of reallocation, and from the elasticity of substitution σ in the market, which governs the welfare gain from a given reallocation. Using a conservative $\sigma = 6$ borrowed from the literature, we find that treating all firms in a market would result in welfare gains of about US\$15,000 per firm, or about 17 percent of profits, largely driven by gains to consumers. Our estimates imply that the private return to capital was about 74 percent per year, most of which was canceled by losses to competitors, and that the social return to capital was about 60 percent, most of which was driven by gains to consumers.

¹In the longer term, if new consumers learn about the improvements in quality and price, the elasticity may increase, and markets may start to grow.

A key implication of these results is that accounting for indirect effects, potentially multiple, can be essential for evaluating firm policies. In our setting, the high social return implies that the program generated large welfare gains. However, accounting for only the direct and indirect effects on firms, while ignoring the effect on consumers, would imply smaller and insignificant gains. Although our returns may seem high, they are comparable to the private returns of 55–63 percent found by De Mel, McKenzie, and Woodruff (2008) and 105 percent found by Banerjee and Duflo (2014). In fact, our analysis may help explain why private returns in developing countries are high. These returns depend on the (quality-adjusted) productivity gain enabled by the loan *times* the potential for business stealing. Thus, even moderate productivity gains can generate high returns when business-stealing effects are strong.

Our research contributes to two main strands of the literature. Our analysis of access to finance builds on work documenting that firms face credit constraints. In particular, De Mel, McKenzie, and Woodruff (2008) and McKenzie and Woodruff (2008) experimentally evaluate the impact of cash grants on microenterprises in Sri Lanka and Mexico, and Banerjee and Duflo (2014) evaluate the impact of a targeted lending program on midsize firms in India.² Our main contributions to this work are the evidence on indirect effects, the mechanisms, and the welfare evaluation.

Our analysis of indirect effects builds on studies of indirect and equilibrium effects in different contexts. Duflo and Saez (2003) introduced the idea of experimental variation in treatment intensity for documenting indirect effects.³ Research on indirect effects involving firms includes Bloom, Schankerman, and Van Reenen (2013), who study the spillover and business-stealing effects of R&D in observational data. Experimental work includes Drexler, Fischer, and Schoar (2014) and Calderon, Cunha, and De Giorgi (2020), who find suggestive evidence of negative indirect effects of financial and business literacy interventions but do not have the design or power to investigate them. Closest to our work, Rotemberg (2019) documents a negative indirect effect on peer firms of a subsidy policy in India, and McKenzie and Puerto (2021) document no indirect effect on peer firms of a microenterprise training experiment in Kenya. Our main contributions to these studies are the randomized evidence on business stealing, the evidence on the mechanisms depicted in Figure 1, and the welfare evaluation. Also related are Burke, Bergquist, and Miguel (2018); Huber (2018); Breza and Kinnan (2021); and Buera, Kaboski, and Shin (2021), who study the (generally positive) general equilibrium impacts of credit, and Sraer and Thesmar (2020), who propose a method to evaluate the general equilibrium effects of firm policies. Our main contributions to their work are the evidence on business stealing, the underlying mechanisms, and the welfare implications.

²Related work exploits shocks and policy variation to measure the impact of bank loans on firm performance, for example, Chodorow-Reich (2014); Ponticelli and Alencar (2016); and Brown and Earle (2017). Also related are studies evaluating the impact of microfinance, including Angelucci, Karlan, and Zinman (2015); Attanasio et al. (2015); Augsburg et al. (2015); Banerjee et al. (2015); Crepon et al. (2015); and Tarozzi, Desai, and Johnson (2015).

³This design has subsequently been used in many domains involving individuals, including financial transfers (Angelucci and Giorgi 2009), labor market policies (Crepon et al. 2013), and the adoption of health products (Guiteras, Levinsohn, and Mobarak 2019).

I. Context, Design, and Data

A. Context and Design

Our experimental site is a province located in southeastern China. The province is a rapidly growing part of China, with average annual GDP growth between 2010 and 2019 of 8.25 percent and GDP per capita in 2021 of over US\$8,000. Our study was conducted in one large prefecture-level city, which consists of several counties and covers over 20,000 square kilometers.

We worked with firms based in the 78 local “markets” in that city. A market is a government-defined geographic cluster of firms mostly specialized in retail, similar to a mall or a bazaar. Markets are located in urban areas and are integrated into the local economy. Each market specializes in a broad product category, such as building materials; most firms in the market sell products in that broad category, but each firm has a more specialized product category (online Appendix Table A1 lists the broad product categories of our markets). For example, in the market for building materials, firms may specialize in selling bricks, wooden flooring, painting materials, or stone. We refer to a firm’s peers in the market that have the same specialized product category as the firm’s competitors. Firms tend to have several competitors in the market. Thus, our setting features a high spatial concentration of competing establishments, a feature that appears to be common in the retail sector in both developing and advanced economies (Leonardi and Moretti 2022).

Markets are far from each other: no county in the city has two markets that are close competitors in terms of their broad product category, and the shortest distance between the centers of any two counties is over 65 kilometers. Each market has a market office and a manager. The duties of the market office include dealing with new applications by firms, allocating firms to the market, collecting rent, gathering data on the activities and performance of firms, monitoring firms to ensure that they make the rent payments, and producing reports about the market for the government.

Partner Bank and Loan Product.—We conducted the intervention in collaboration with our partner, a major commercial bank in China. In 2013, the bank introduced a new loan product to SMEs in the study province. An important feature was that the loan product was introduced to firms in local markets of the type just described.

The new loan product had several characteristics potentially attractive to borrowers.⁴ Most importantly, it did not require collateral, whereas most alternative formal loans seemed to require either collateral or a guarantor who was a government employee. In addition, the new loan had a standardized application form, and the bank committed to make a decision in two weeks. Finally, the monthly interest rate was about 0.7 percent, which was at a 15 percent discount relative to existing formal loans in our baseline survey.

⁴In our baseline survey, we asked nonborrowing firms why they did not have a loan. Thirty-one percent reported that they did not need credit. Among the rest, 40 percent said that they did not possess collateral, 39.6 percent that the time cost of applying was too high, 22 percent that they could not borrow enough, 14.8 percent that the interest rate was too high, and 11.6 percent that they did not know how to apply. Note that we allowed respondents to choose multiple answers.

A natural question is why the bank was able to offer these loan terms, which were meaningfully better than those of competitors. Our interactions with bank officials suggest that one main reason was that the firms were organized in markets, which allowed the bank to economize on the screening and monitoring costs that Banerjee and Duflo (2010) identify as a key barrier to lending. In particular, loan officers could gather information from the market manager about potential borrowers and could save on transportation costs by visiting multiple borrowers in the same market. These issues seemed to matter for the bank, which engaged in intensive screening and monitoring activities, including several visits to borrower firms before and after approving the loan. A second reason may have been that there were dynamic incentives: because no other bank offered uncollateralized loans, a firm that lacked collateral would have lost its ability to borrow if it did not repay the loan. Consistent with these reasons, the approval rate for loans in our sample was only about 47 percent, but the repayment rate for approved loans was about 98 percent.

The loan worked as follows. Firms interested in the loan could apply irrespective of their treatment status, and the bank would screen them using their screening process. Once a borrower was approved, the bank assigned them a credit limit of up to 30 percent of the value of net assets as computed by the bank, capped at a maximum loan amount of ¥500,000 (about US\$81,000). Firms could borrow any amount up to their assigned credit limit. They then had to make monthly interest payments, and to repay the loan within two years. Taking out a new loan after repayment was possible.

Intervention.—In the spring of 2013, we obtained from the market offices the lists of all active firms in each market. The lists included the number of formal employees of each firm. The total number of firms in the 78 markets was over 6,000. In the summer of 2013, after the baseline survey (described below), we introduced our intervention to this population of firms.

Our intervention combined a market-level and a firm-level randomization. Out of our 78 markets, we randomized 37 markets to have high treatment intensity, 10 markets to have medium treatment intensity, and 31 markets to have zero treatment intensity (pure control).⁵ In high-treatment-intensity markets, we treated a randomly selected 80 percent of firms; in medium-intensity markets, we treated a randomly selected 50 percent of firms; and in pure control markets, we treated no firms. In the market-level randomization, we stratified markets by county and, within each county, by whether the number of firms in the market was above or below the median in the county, resulting in 22 strata. In the firm-level randomization, we stratified firms by whether the number of employees, as provided by the market office, was above or below the median in the market.

The specific treatment was as follows. Every treated firm was visited every month for a year by a loan officer who provided information about the new loan product and, if the firm decided to apply, provided help with the application process, including filling in the relevant forms. Once a treated firm decided to borrow, the loan officer stopped visiting. In addition, the bank agreed not to send loan officers

⁵Our target was to have 45 percent of markets in the high-intensity treatment, 15 percent of markets in the medium-intensity treatment, and 40 percent of markets in the pure control.

to visit untreated firms for a year. However, untreated firms could still learn about the loan product from other sources, including their social network, and were free to apply for a loan.⁶ Once a firm—whether treated or untreated—submitted an application, the bank decided on lending independently of us.

Surveys.—Because the total number of firms, over 6,000, was beyond our capacity to survey, we randomly selected half of the firms, in each strata of each market, to be included in our survey sample. This gave us 3,173 firms. We conducted three long surveys with this sample of firms: a baseline survey in the summer of 2013, before the intervention; a midline survey in the summer of 2015, two years after the intervention; and an endline survey in the summer of 2016. Because the fiscal year in China ends in June, data in the baseline survey refer to the fiscal year before the intervention. At the baseline survey, treated firms did not yet know that they would soon be receiving a treatment. We waited two years between the baseline and the midline surveys to allow time for the firms to borrow and use the loan. In addition to these long surveys, we conducted a short follow-up survey in the summer of 2020, in which we collected some additional information not covered in our long surveys.

The surveys were conducted by locally hired enumerators in collaboration with the market office and with the bank. A member of the market office guided us and introduced us to the manager of each firm. At the introduction, a loan officer from the bank was also present to ensure that the manager would consider us trustworthy. The market and bank officials then left, and the survey was conducted, in person, with the manager of the firm.

In the long surveys, we collected information about the following groups of variables. (i) Firm characteristics, including sales, profits, employment, various cost categories, and other balance sheet variables. We collected two measures of sales. Besides the self-reported value, which we asked for in all three long surveys, in the endline survey we also collected the book value directly by having the enumerator ask the firm's accountant or manager to physically show the value in the firm's books. (ii) Managerial characteristics, including demographics. (iii) Measures of financial and business activities, including borrowing from formal and informal sources, the use of trade credit, the number of suppliers and clients, measures of product introduction, renovation, advertising, and others. In addition, in the endline survey we also asked borrowers what they primarily used the loan for.

The 2020 short follow-up survey had three components: a market survey, a firm survey, and a consumer survey. In the market survey, we collected information from the market office about the four closest neighbors of each firm. In the firm survey, we asked retrospective questions about outcomes that we did not include in the long surveys, especially the average price of their main product in 2016 and the share of employees with at least high school education in 2016. In the consumer survey, we randomly picked a customer visiting the firm, and asked her or him a number of questions about product and service quality. Our questions included whether the customer received valuable advice from the seller, and their evaluation

⁶In principle, untreated firms could also learn about the loan by observing the loan officer visiting a treated firm. Because firms have many customers, and because loan officers do not wear uniforms, this mechanism was unlikely to have been important.

on a five-point scale of the shopping environment, service quality, value for money, and overall satisfaction.

Borrowing Data from the Bank.—We complement the survey data with data from our partner bank. These data include which of the firms in our sample borrowed using the new loan product, the month they borrowed, the interest rate, and the loan amount.

B. Summary Statistics

Our full sample consists of 3,173 firms organized in 78 markets. In the average market, we observe about 41 firms, which specialize in about 4.9 main products, so that there are 4.9 distinct groups of competitors. Since we only surveyed half of the firms in each market, it follows that the average market had about 82 firms specializing in about 5 main products, with around 17 firms specializing in each main product. This confirms that firms operate in a fairly competitive environment.

Table 1 presents summary statistics in the baseline survey about firms and managers. Each row corresponds to a separate regression, in which the variable listed in the first column is regressed on a constant and on four indicators for the different treatment arms: treated firm in a 50 percent market, untreated firm in a 50 percent market, treated firm in an 80 percent market, and untreated firm in an 80 percent market. Thus, the coefficient of the constant measures the average of the variable in the “pure control” group of firms in untreated markets, while the other coefficients measure the average difference in that variable between the other treatment arms and the pure control. We label the columns accordingly. We cluster standard errors at the market level.

Panel A on firm characteristics shows that in 2013 average firm age was about 6.5 years. Almost 70 percent of the firms were in retail, with the rest mostly in services and manufacturing. Firms employed on average about nine workers. The average net profit was about ¥519,000 (about US\$84,000), and average revenue was about ¥3,230,000 (US\$525,000). Panel B presents managerial characteristics. Almost 60 percent of managers were men, and in 2013 they were on average 38 years old. About a quarter of them had a college degree. About 15 percent of managers had a political connection, defined as past experience working in the government. Panel C reports on borrowing from formal banks. A quarter of firms had a preexisting loan at baseline. Conditional on having a loan, the average loan amount was about ¥300,000 (US\$49,000), and the average monthly interest rate was about 0.9 percent. Panel D reports data on business connections with suppliers and clients. Firms had about 27 clients per day and about 6 active suppliers. Consistent with the randomization, there is no significant difference between any of the treatment arms in any of these variables. Moreover, as the p -values at the bottom show, regressing the treatment indicators on all of the variables in panels A–D yields no evidence of joint significance.

Panel E reports attrition and shutdown by endline. Attrition is defined as one in a survey wave if we do not have information about the firm in that wave and do not know whether the firm has shut down; this is typically due to the firm moving away or choosing not to respond. Shutdown is defined as one in a survey wave if we have

TABLE 1—SUMMARY STATISTICS AND BALANCE

Sample: All baseline, 3,173 firms	Pure control	Δ Treated 50% markets	Δ Untreated 50% markets	Δ Treated 80% markets	Δ Untreated 80% markets
<i>Panel A. Firm characteristics</i>					
Firm age	6.479 (0.308)	0.778 (1.007)	1.009 (0.724)	-0.214 (0.420)	-0.447 (0.469)
Sector—retail (%)	0.682 (0.057)	0.047 (0.089)	0.027 (0.103)	0.004 (0.072)	-0.041 (0.090)
Number of employees	8.823 (0.564)	1.159 (1.151)	0.364 (1.131)	0.015 (0.705)	0.219 (0.697)
Profit (¥10,000)	51.95 (6.193)	-1.878 (11.62)	-2.483 (9.134)	-0.951 (7.747)	-0.272 (8.204)
Sales (¥10,000)	323.7 (38.30)	19.06 (79.75)	6.570 (59.83)	2.925 (53.74)	-7.416 (43.40)
<i>Panel B. Managerial characteristics</i>					
Gender (1 = Male, 0 = Female)	0.581 (0.031)	-0.018 (0.065)	-0.009 (0.061)	-0.002 (0.053)	-0.002 (0.059)
Age	38.36 (0.642)	-0.232 (1.415)	0.347 (1.294)	-0.016 (1.081)	0.927 (1.059)
Education—college	0.246 (0.021)	0.011 (0.036)	0.025 (0.051)	0.031 (0.028)	0.029 (0.034)
Political connection (1 = Yes, 0 = No)	0.148 (0.018)	0.037 (0.0400)	0.015 (0.031)	0.015 (0.025)	0.013 (0.027)
<i>Panel C. Borrowing</i>					
Other bank loan (1 = Yes, 0 = No)	0.253 (0.024)	0.036 (0.049)	-0.001 (0.048)	-0.027 (0.033)	-0.030 (0.044)
Loan size (¥10,000)	30.78 (6.737)	1.271 (14.28)	-4.008 (8.919)	-1.982 (11.12)	-5.531 (7.769)
Monthly interest rate (%)	9.158 (0.133)	-0.391 (0.380)	0.332 (0.289)	0.043 (0.198)	0.036 (0.294)
<i>Panel D. Partnerships</i>					
Number of clients	27.37 (1.011)	-0.770 (1.505)	1.232 (2.287)	1.124 (1.482)	2.118 (1.829)
Number of suppliers	6.535 (0.813)	2.091 (2.245)	1.549 (1.559)	-0.244 (0.908)	0.124 (1.063)
<i>Panel E. Attrition and shutdown (endline)</i>					
Attrition	0.106 (0.009)	-0.002 (0.015)	-0.002 (0.023)	0.001 (0.012)	-0.001 (0.016)
Shutdown	0.134 (0.023)	-0.026 (0.059)	-0.031 (0.045)	-0.052 (0.028)	0.019 (0.034)
<i>p</i> -val of joint significance of vars in panels A–D		0.439	0.354	0.716	0.328
Observations	1,247	222	203	1,214	287

Notes: Each row is a separate regression of the outcome variable (column 1) on a constant representing the pure control group, and indicators for treated firms in 50 percent markets, untreated firms in 50 percent markets, treated firms in 80 percent markets, and untreated firms in 80 percent markets, representing the mean differences relative to the pure control. Panels A–D focus on balance at baseline, panel E on attrition and shutdown at endline. To test for joint balance, we regress the four treatment indicators on all baseline variables of panels A–D and report the *p*-value of joint significance. Standard errors clustered at the market level.

information that the firm went out of business in or before the fiscal year to which the survey wave refers. With these definitions, attrition and shutdown are mutually exclusive. We made several arrangements to keep attrition low. The bank phoned managers in advance to arrange the survey; when the manager was unavailable at the

arranged time, we attempted to arrange a second meeting, and with the help of the bank and the market office, we managed to track most mover firms. The table shows that attrition was about 10 percent by endline and not significantly different across treatment arms. Finally, the shutdown rate among pure control firms was about 13 percent by endline and was significantly lower ($p < 0.10$) among treated firms in the 80 percent treated markets. This result suggests that the treatment may have improved firm performance but also that differences in shutdown rates may have induced selection by treatment status, potentially biasing our results. To address this concern, in online Appendix Table A2, we present balance tests for the subsample of firms that remain in our data all the way to the 2016 endline, or to the 2020 follow-up, and document no significant differences by treatment status in these subsamples. These results suggest that selective exit is unlikely to bias our results.

II. Model and Empirical Strategy

To motivate our empirical analysis, in this section we build a model of the direct and indirect effects of the loan program. We begin with a model of business stealing, then incorporate take-up and information diffusion, and then use the predictions to formulate our empirical strategy.

A. A Model of Business Stealing

Our basic model parallels that of Rotemberg (2019) but explicitly incorporates markets and the possibility that the loan enables quality improvements. The main focus of our model is a monopolistically competitive sector, which consists of a mass of markets indexed by m , and in each market m a mass n_m of firms indexed by i . Goods purchased in market m aggregate into a composite good

$$(1) \quad Q_m = \left[\int_{i \in m} (h_i \cdot Q_i)^{1-1/\sigma} di \right]^{\frac{\sigma}{\sigma-1}},$$

where h_i is the quality or appeal of the product (or service) of firm i . Consumer preferences are given by

$$(2) \quad H + \left(\int Q_m^{1-1/\theta} dm \right).$$

In these equations, σ measures the elasticity of substitution within a market, and θ the elasticity of substitution across markets. We assume $\sigma > \theta > 1$. For simplicity, we assume that utility is quasi-linear, with H being a numeraire good produced by a perfectly competitive sector and traded at a price normalized to one. One unit of labor produces w units of H , pinning down the wage as w . The aggregate labor supply is L .⁷ Firms in the monopolistically competitive sector have constant returns to scale, produce with labor, and take wages as given. The output of firm i is $Q_i = \omega_i L_i$. Firms may differ both in quality (or appeal) h_i and productivity ω_i , and different markets m may have different distributions of firm quality and productivity.

⁷We assume that L is sufficiently large to ensure that in equilibrium $H > 0$.

Recognizing quality or appeal is important, given the finding of Hottman, Redding, and Weinstein (2016) that in the United States it accounts for most of the variation in retail firm size. We note that we interpret quality in a broad sense that also includes the product variety offered by the firm.⁸

We consider the impact in this economy of a loan program that provides loans to a subset of firms in a subset of markets. For now, we take the loan assignment to be exogenous and assume that all firms assigned the loan borrow; we will endogenize the borrowing decision below. We assume that receiving the loan allows for investments in business practices that improve both product quality and firm productivity: quality h_i increases by a factor $e^{\gamma h}$, and productivity ω_i by a factor $e^{\gamma \omega}$. We assume $\gamma_\omega, \gamma_h \geq 0$, and let $\gamma_\omega + \gamma_h = \gamma$. We let B_i be an indicator for whether firm i borrows and let Z_m denote the share of firms in market m that borrow.

In the spirit of the potential outcomes approach, it will be helpful to consider counterfactual outcomes that would obtain absent the loan program. We use the convention that variables with tilde represent outcomes absent the loan program, and Δ represents the impact of the program, i.e., $\Delta X = X - \tilde{X}$.

Because the loan assignment need not be random, borrowers may be different from the average firm in their market. To capture this selection, we let λ_m denote the ratio, absent the loan program, of the average revenue of borrowers relative to that of all firms in market m :

$$\lambda_m = \frac{\int_{i \in m: B_i=1} \tilde{R}_i di}{Z_m \cdot \int_{i \in m} \tilde{R}_i di}$$

where \tilde{R}_i is the revenue of firm i absent the intervention. Then the following result characterizes the impact of the loan program on firms.

PROPOSITION 1: *To a first-order approximation, the impact of the loan program on the revenue of firm i in market m is*

$$(3) \quad \Delta \log R_i \approx (\sigma - 1)\gamma \cdot B_i - (\sigma - \theta)\gamma \lambda_m \cdot Z_m.$$

All proofs are in online Appendix A.1. The impact on revenue is characterized by two terms. The first term represents the positive direct effect of the firm receiving a loan (B_i), while the second term represents the negative business-stealing effect of the share of the firm’s competitors that receive the loan (Z_m). The logic of the direct effect is that the loan induces (i) improvements in quality (and variety) and (ii) reductions in cost, both of which allow the firm to reduce its quality-adjusted price and thereby attract higher demand. These mechanisms are depicted in Figure 1. The coefficient is $\gamma(\sigma - 1)$, where γ measures the decrease in the quality-adjusted price and $\sigma - 1$ measures the demand response.

The logic of the business-stealing effect is that an increase in the share of competitors who borrow increases market-level competition (formally, reduces the quality-adjusted price index) and thus lowers demand for the product of the firm.

⁸ It is straightforward to explicitly incorporate product variety in the model, and improving product variety is equivalent to improving the firm’s quality-to-cost ratio.

This is the flip side of the direct effect. The coefficient is proportional to λ_m because larger borrowers have a larger absolute impact on the market. The coefficient is also proportional to $\sigma - \theta$, which measures the responsiveness of firm-level demand to market-level competition. This elasticity is lower than the $\sigma - 1$ that governs the direct effect because an increase in market-level competition also attracts demand from outside of the market.⁹

Observe that with $\lambda_m = 1$, in the limiting case when θ approaches one, the direct and indirect effects become approximately equal. This case corresponds to Cobb-Douglas preferences over markets, which imply that the relative expenditure share of a market does not respond to market-level improvements. Importantly, this market-level nonresponse can be consistent with substantial within-market reallocation as characterized by the proposition.

B. Diffusion, Take-Up, and Other Empirically Relevant Features

We enrich the model by incorporating the indirect effect of information diffusion, as well as other features relevant for our empirical analysis: multiple periods, random shocks, the randomized intervention, and imperfect take-up.

We assume that the model is repeated over periods $t = 0, 1, \dots, \tau$. Consumers have preferences given by (2) in each period and consume all their income each period. We make relatively weak assumptions about firm dynamics, stated formally in the online Appendix. Specifically, absent the treatment, for each firm i , the vector of log quality and productivity $\tilde{x}_i^t = (\log \tilde{h}_i^t, \log \tilde{\omega}_i^t)$ evolves according to the sum of (i) a Markov process that depends on firm-level and market-level characteristics and shocks, (ii) a firm-specific idiosyncratic shock, and (iii) a time trend. We also assume that wages w^t , pinned down by the production function for the numeraire good, evolve deterministically.

We model the intervention as follows. A mass M of markets are selected to be treated. Market-level treatment intensities are characterized by intensity levels $0 \leq s_1 \leq s_2 \leq \dots \leq s_k$ and associated probabilities $0 \leq q_j \leq 1$ such that $q_1 + \dots + q_k = 1$. The intervention is introduced between periods 0 and 1 and treats a randomly chosen share s_j of firms in a randomly chosen share q_j of markets.

The treatment provides information about the new loan opportunity. Similar to the basic model, we assume that every firm that takes up the loan experiences increases in (log) quality and productivity of γ_h and γ_ω , and we also assume that these improvements are permanent. But, different from the basic model, information about the loan may diffuse to untreated firms, and both treated and untreated firms make a decision about whether to borrow.

We model information diffusion in the spirit of the Bass (1969) model by assuming that, when a share S_m of firms in market m are treated, an additional share $\phi S_m(1 - S_m)$ learn about the loan opportunity. This share is proportional to the mass of treated firms S_m that can potentially diffuse the information and the mass of untreated firms $1 - S_m$ that can potentially receive the information. The parameter ϕ governs the strength of diffusion. We assume that the firms to which

⁹In our model with quasi-linear preferences, the reallocation from outside of the market is coming entirely from the numeraire good.

information diffuses are a random subset of the untreated firms. These assumptions imply that the treatment only generates information diffusion about the loan opportunity. However, the model permits non-treatment-induced information diffusion about quality and productivity by allowing these variables to have a market-level component.

We assume that take-up is imperfect because the firm’s manager may not have a sufficiently promising idea that could be developed using the loan. Furthermore, we assume that, relative to the benefits of the loan, the cost of applying is minimal and can be ignored. Treated firms decide to take up in period 1, while informed untreated firms, because diffusion takes time, decide to take up in an exogenously given period s , where $s \geq 1$. We model take-up in a reduced-form fashion that allows for a correlation between firm fundamentals and whether the manager has an idea for the loan. Specifically, firm i , if treated, takes up at $t = 1$ with probability $F^T(\tilde{x}_i^1)$, and if untreated but reached by diffusion, takes up at $t = s$ with probability $F^D(\tilde{x}_i^s)$. Here, the exogenous nondecreasing functions F^T and F^D represent the probability that the firm with fundamentals $\tilde{x}_i^t = (\log \tilde{h}_i^t, \log \tilde{\omega}_i^t)$ has an idea.¹⁰

Denote the probability that in market m a random treated firm borrows by $\mu_m^T = E[F^T(\tilde{x}_i^1) | m]$, and that in market m a random firm accessed by diffusion borrows by $\mu_m^D = E[F^D(\tilde{x}_i^s) | m]$. We assume that $\mu_m^T > \phi \mu_m^D$ for all m , which means that, on average, getting the treatment has a higher effect on take-up than getting diffusion from all peers. This assumption will be useful for our IV strategy. Let μ^T and μ^D denote the unconditional average take-up probability of a treated and a diffusion firm.

We let T_i denote the realized treatment status of firm i , and S_m the realized treatment intensity of market m . Because borrowing is now time-dependent, we let B_i^t indicate whether the firm has borrowed in or before period t , Z_m^t denote the share of firms in the market that have borrowed in or before period t , and $\lambda_m^t = \int_{i \in m: B_i^t=1} \tilde{R}_i^t di / (Z_m^t \cdot \int_{i \in m} \tilde{R}_i^t di)$ denote, absent the treatment, the average revenue of borrowers relative to the average revenue of all firms in market m in period t .

PROPOSITION 2:

1. Borrowing in a period $t \geq s$ can be written as

$$(4) \quad B_i^t = \mu^T \cdot T_i + \mu^D \phi \cdot (1 - T_i)S_m + \eta_i,$$

where

$$E[\eta_i \times (T_i, (1 - T_i)S_m)] = 0.$$

2. Firm revenue in period t , to a first-order approximation, can be written as

$$(5) \quad \log R_i^t \approx (\sigma - 1)\gamma \cdot B_i^t - (\sigma - \theta)\gamma \lambda \cdot Z_m^t + \kappa \cdot Post^t + f_i + \varepsilon_i^t,$$

¹⁰ At the expense of additional notation, we could allow take-up to also depend on market-level characteristics.

where λ is a non-negative-weighted average of λ_m^t , $Post^t$ is an indicator for $t \geq 1$, f_i are firm-specific effects, and $E[\varepsilon_i^t \times (1_j, Post^t, Post^t \cdot T_i, Post^t \cdot S_m)] = 0$ for all j .

Part 1 characterizes take-up. The first term on the right-hand side of (4) shows that, on average, receiving the treatment increases the probability of borrowing by μ^T . The second term characterizes diffusion: a nontreated firm ($1 - T_i = 1$) in market m is reached by information diffusion with probability ϕS_m and conditional on being informed, borrows with average intensity μ^D . The error term η_i reflects both firm-level and market-level idiosyncratic variation in take-up and is orthogonal to the explanatory variables because of the random treatment assignment.

Part 2 characterizes firm revenue. Equation (5) parallels equation (3) of Proposition 1 but expresses revenue, rather than the treatment effect, and incorporates the additional richness of the model. The direct effect, $(\sigma - 1)\gamma \cdot B_i^t$, is the same as in the previous result, except that it now accounts for time variation in borrowing. In the business-stealing effect, $-(\sigma - \theta)\gamma\lambda \cdot Z_m^t$, a novelty relative to the previous result is that the coefficient λ that accounts for selection is no longer dependent on m (or on t). To obtain this term, using a logic familiar from the study of heterogeneous treatment effects, we move the heterogeneity in the business-stealing effect captured by λ_m^t to the error term. The coefficient of the business-stealing effect thus represents an average effect and depends on a weighted average λ , which measures average selection.¹¹ The equation also includes firm and time fixed effects that account for other heterogeneities and dynamics.

The main novelty in (5) is the error term, which captures both firm-level shocks and—as we have just seen—heterogeneity in the strength of the business-stealing effect. Our model allows both of these sources of variation to be correlated with take-up and hence the right-hand-side variables. Thus, equation (5) cannot be estimated with ordinary least squares. However, the proposition states that—essentially because the treatments are randomly assigned—the error term is orthogonal to $Post^t \cdot T_i$ and $Post^t \cdot S_m$. This implies that equation (5) can be estimated with an IV, using $Post^t \cdot T_i$ and $Post^t \cdot S_m$ as instruments.

The reduced-form equation of this IV strategy takes the familiar difference-in-differences form

$$(6) \quad \log R_i^t = \beta \cdot T_i \cdot Post^t + \delta \cdot S_m \cdot Post^t + \kappa \cdot Post^t + f_i + \epsilon_i^t.$$

In the important special case of no diffusion, we can use the take-up equation (4) to substitute borrowing outcomes with the treatments, resulting in an explicit expression for the coefficients of this reduced-form regression: $\beta = (\sigma - 1)\gamma \cdot \mu^T$ and $\delta = (\sigma - \theta)\gamma\lambda \cdot \mu^T$. Thus, in this special case the reduced-form coefficients are proportional to the second-stage coefficients and serve as measures of the direct and business-stealing effects.

¹¹In this average effect, the weights are functions of the covariance between the share of firms treated and the share of firms that borrow. The nonnegativity of this covariance, and hence the weights, is ensured by our assumption that take-up responds more to the treatment than to diffusion.

C. Empirical Strategy

Proposition 2 motivates our empirical strategy. Part 1 of the proposition provides foundations for our first estimating equation, which documents information diffusion in take-up:

$$(7) \quad \text{Borrow}_i = \text{const} + \mu \cdot \text{Treated}_i + (1 - \text{Treated}_i) \\ \cdot G(\text{SharePeersTreated}_i) + \eta_i.$$

Here, the $G(\cdot)$ function governs how diffusion varies with the treatment intensity and is assumed to be linear in the model. In our main specification, we define the set of peers relevant for diffusion to be all firms in the market, but later, we also consider alternative definitions based on distance and competition status.

Our second estimating equation documents indirect effects in firm performance:

$$(8) \quad y_i^t = \beta \cdot \text{Post}^t \times \text{Treated}_i + \delta \cdot \text{Post}^t \times \text{ShareCompetitorsTreated}_i \\ + \kappa \text{Post}^t + \text{Firm f.e.} + \varepsilon_i^t.$$

When the outcome variable y_i^t is log revenue, this specification is simply the reduced-form (6) of the IV in Proposition 2. As discussed above, within our model, in the case of no diffusion, the reduced-form coefficients have clear interpretation: β is proportional to the direct effect, and δ is proportional to the business-stealing effect of the loan. With diffusion, the map from reduced-form coefficients to economic forces is more complicated. However, even then, $\beta > 0$ and $\delta < 0$ are only possible if the loan has a positive direct and a negative business-stealing effect.¹² Moreover, independent of the model, the coefficients have straightforward interpretations as treatment effect estimates: β measures the direct and δ the indirect effect of the treatment. For these reasons, we use (8) as our main estimating equation. In our main specifications, we define competitors to be firms that are in the same market and have the same specialized product category.

A key identification condition of equation (8) is that the *ShareCompetitorsTreated* variable is as good as random. In our empirical setting this condition may not seem immediate since the share of competitors treated is not a treatment arm. But note that the identification condition follows mechanically if each competitor has the same exogenous probability of receiving the treatment. Indeed, in that case, the expected share of competitors treated, conditional on any firm-level disturbance, is simply that exogenous probability, implying mean independence. In our empirical setting—except for the integer constraints of the randomization—each firm has the same probability of receiving the treatment. Thus, the identification condition

¹²Diffusion biases β toward zero because some untreated firms also borrow and experience the direct effect. Thus, $\beta > 0$ must imply a positive direct effect. Given that, $\delta < 0$ must imply a negative indirect effect since otherwise, diffusion would just amplify the positive direct effect.

should hold, and balance tests in online Appendix Table A3 confirm that the *ShareCompetitorsTreated* is uncorrelated with baseline characteristics.¹³

Our third estimating equation jointly accounts for the diffusion and business-stealing effects through the IV regression

$$(9) \quad y_i^t = \zeta \cdot Borrow_i^t + \xi \cdot Share\ Competitors\ Borrow_i^t + \kappa Post^t \\ + Firm\ f.e. + \nu_i^t,$$

in which $Post^t \times Treated_i$ and $Post^t \times ShareCompetitorsTreated_i$ are the instruments. When y_i^t is log revenue, this specification is identical to the IV of Proposition 2, providing a way to infer, even with diffusion, the direct and business-stealing effects of the loan. We use this regression to confirm the insights obtained from the reduced-form analysis, and for our welfare evaluation.

Before moving to the results, we explain how our approach can account for two different indirect effects using a single source of exogenous variation: treatment intensity. The reason is that the two indirect effects affect different outcome variables in different regressions: diffusion affects take-up in (7), while business stealing affects firm performance conditional on take-up in (9).

III. Reduced-Form Estimates of Direct and Indirect Effects

In this section, we present five sets of reduced-form results. First, we report impacts on loan take-up. Second, we report impacts on several firm performance measures and examine the robustness of these results. Third, we look at intermediate outcomes to document evidence on mechanisms. Fourth, we report impacts on market-level outcomes. Finally, we explore the heterogeneity of the indirect effects with respect to distance and competition status, and in the process identify a third indirect effect: diffusion of demand.

A. Take-Up

Table 2 presents cross-sectional estimates of our first estimating equation (7). In columns 1–3, the dependent variable is an indicator for whether the firm has borrowed using the new loan product by the endline survey. Column 1 shows that the probability of borrowing was 28 percentage points higher among treated firms than among untreated firms, indicating that the treatment succeeded in inducing firms to borrow. Complementing this finding, the data show that the average loan amount borrowed using the new product was ¥270,000 (about US\$47,000), or roughly 9 percent of average sales. The average monthly interest rate was about 0.73 percent,

¹³Borusyak and Hull (2022) show that in cases when the indirect effect contains nonrandom variation driven by a variable mediating exposure—e.g., when it is governed by the number of competitors treated—one should “re-center” the indirect effect by subtracting its mean conditional on the mediating variable. When, as in our case, the indirect effect is a share, that conditional mean is a constant, and re-centering does not make a difference. Aronow and Samii (2017) study the case where the indirect effects associated with different treatment intensities are not linked. Our model implies that they are a linear function of the intensity level, a prediction on which we present evidence below.

TABLE 2—EFFECTS ON BORROWING BY ENDLINE

	Borrow with new loan product			Borrow from other sources		Borrow from any source	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated	0.279 (0.034)	0.315 (0.034)		0.029 (0.019)		0.274 (0.034)	
Untreated × Share of Peers Treated		0.178 (0.037)		0.013 (0.032)		0.164 (0.041)	
Treated × 50% Market			0.302 (0.057)	0.029 (0.025)		0.280 (0.049)	
Treated × 80% Market			0.318 (0.039)	0.028 (0.021)		0.273 (0.038)	
Untreated × 50% Market			0.112 (0.062)	0.005 (0.028)		0.096 (0.059)	
Untreated × 80% Market			0.140 (0.030)	0.007 (0.029)		0.137 (0.033)	
Constant	0.067 (0.014)	0.032 (0.037)	0.031 (0.030)	0.294 (0.013)	0.295 (0.013)	0.320 (0.016)	0.322 (0.016)
Mean of pure control Observations	0.031 3,173	0.031 3,173	0.031 3,173	0.295 2,658	0.295 2,658	0.322 2,658	0.322 2,658

Notes: Dependent variable is indicator for borrowing by endline, in columns 1–3 from new loan product; in columns 4–5 from other sources; and in columns 6–7 from any source. Sample in columns 1–3 is all firms, and in columns 4–7 is firms in the endline survey. Share of peers treated is the share of other firms in the market treated. Standard errors clustered at the market level.

meaningfully lower than the average interest rate among alternatives reported in Table 1. And, although the bank did not share data on this with us, they reported that the repayment rate among borrowers in our sample was over 98 percent and that about half of borrowers borrowed again after repaying the loan. We conclude that firms found the new loan product attractive and that the treatment succeeded in providing access to a significant amount of financing.

In columns 2 and 3, we look in more depth at borrowing by untreated firms. Column 2 includes the interaction between an indicator for the firm being untreated and the share of other firms in the market that are treated. The estimated coefficient of 0.18 is highly significant and implies that increasing the share of peers treated from 0 percent to 100 percent would increase the likelihood that an untreated firm borrows by 18 percentage points. Column 3 is a more flexible specification that estimates take-up separately in the five treatment arms. Firms in pure control markets borrowed with a 3 percent probability. Relative to these firms, treated firms in the two types of treated markets borrowed with 30 and 32 percentage points higher probability, respectively. And, relative to pure control firms, untreated firms in the two types of treated markets borrowed with 11 and 14 percentage points higher probability, respectively.¹⁴

By showing that untreated firms in treated markets were more likely to borrow than untreated firms in control markets, these results provide clear evidence of an indirect effect on take-up. One plausible explanation for this effect, highlighted by our model, is information diffusion about the loan. Another explanation may be

¹⁴The linear specification of column 2 would imply 9 and 14 percentage points for the last two effects. Panel A of Figure 3 compares graphically these two specifications and shows that we cannot reject linearity.

cost, if loan officers' more frequent presence in treated markets reduced the cost of application for untreated firms. In our context, the contribution of the cost explanation is probably small. Since loan officers did not wear uniforms, untreated firms could not easily identify them, and, even if they could, loan officers were instructed not to provide the same application assistance to these firms. Responses to our survey questions also support the information diffusion explanation: among untreated borrowers in treated markets, 70 percent heard about the program from other firms, 16 percent from friends or relatives, and 5 percent from bank officers (the remaining categories were media and other). For these reasons, we interpret the indirect effect on borrowing as information diffusion. This interpretation is consistent with evidence from other contexts documenting information diffusion about financial products (Banerjee, Chandrasekhar, Duflo 2013; Cai, de Janvry, and Sadoulet 2015; Cai and Szeidl 2018).

The direct and indirect effects on take-up may reflect either increased total borrowing or the new loan crowding out other borrowing (Banerjee et al. 2015). To distinguish between these possibilities, columns 4 and 5 of Table 2 report regressions analogous to columns 2 and 3 but with an indicator for (formal or informal) borrowing from other sources as the outcome. Both the direct and indirect effects are insignificant and small, implying that the treatment did not crowd out other borrowing. In columns 6 and 7, we report analogous regressions in which the outcome variable is borrowing from any source, including the new loan product. The results are similar to those in columns 2 and 3 and confirm that the loan product created new borrowing both directly and indirectly.¹⁵ Because non-credit-constrained firms, if they borrowed with the new loan product, should have used the proceeds to reduce other borrowing (Banerjee and Duflo 2014), the results also provide evidence that firms in our sample were credit constrained. Analogous regressions for the loan amount, shown in online Appendix Table A4, yield parallel results.

B. Main Effects on Firm Performance

Graphical Evidence.—We start by studying the intervention's impacts on firm performance with graphical evidence. The left panel in Figure 2 plots the kernel density of log sales at baseline for three different groups of firms: treated firms, untreated firms in treated markets, and untreated firms in control markets. Consistent with the randomization, these densities are close to each other: before the intervention, the distribution of log sales was similar in the three groups.

The right panel of Figure 2 shows, for the same three groups, the kernel density of the *change* in log sales between baseline and endline. There are two salient differences between these densities. First, relative to untreated firms in control markets, treated firms experienced higher growth in sales. This pattern suggests that the intervention had a positive direct effect on the revenue of treated firms; presumably, access to the loan allowed firms to change business practices and expand. Second, relative to untreated firms in control markets, untreated firms in treated markets

¹⁵The sample in columns 4–7 is smaller than that in columns 2–3. This is because data on other borrowing are from the survey and only available for firms surveyed at the endline, while the data on the new loan are from the bank and available for all firms. Rerunning columns 2–3 in the subsample of columns 6–7 yields similar results.

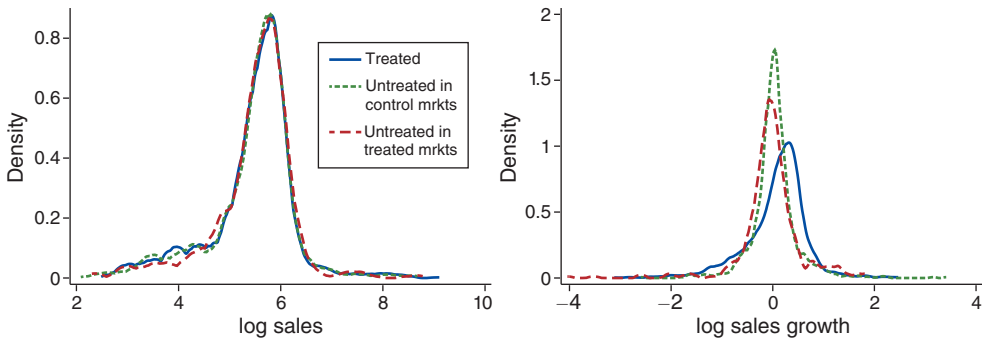


FIGURE 2. KERNEL DENSITY OF LOG SALES

experienced *lower* growth in sales. This pattern suggests that the intervention had a negative indirect effect on the revenue of untreated firms that had treated peers. A natural interpretation is business stealing: as treated firms expanded due to the loan, they lured away some of the business of their untreated competitors.

Regression Estimates.—We now study these direct and indirect effects in more depth using the empirical strategy implied by the model. Table 3 reports estimates of our reduced-form estimating equation (8) for a range of leading firm performance measures. In column 1, the outcome is log sales. The coefficient of the interaction between *Treated* and *Post* implies that the treatment would increase firm sales by 9.9 log points. The coefficient of the interaction between *ShareCompetitorsTreated* and *Post* implies that treating all of its competitors would reduce firm sales by 8.6 log points. Both of these effects are large and highly significant. In column 2, the outcome is profit in levels: because profit may be negative, we do not take logs. We estimate a positive and significant direct effect of ¥126,400 (about US\$20,000) and a negative and significant indirect effect of ¥95,000 (about US\$15,000).

These results, consistent with the logic of the model, suggest that the intervention induced a reallocation of demand and profit from firms that had many treated competitors to treated firms. The similar magnitude of the direct and indirect effects suggests that the treatment had little aggregate effect on market-level revenue or profit. Observe that the latter conclusion holds for any level of treatment intensity because treatment intensity scales both the direct and the indirect effect: the former by increasing the share of firms that receive the treatment and the latter by increasing the strength of the business-stealing effect.¹⁶

The remaining columns in Table 3 focus on factor use, input use, and firm survival. Columns 3 and 4 show positive direct and negative indirect effects for log employment and the log wage bill. These results suggest that employment was an important margin of adjustment accommodating the reallocation induced by the

¹⁶For example, suppose that 50 percent of firms are treated. Then the aggregate direct effect is $9.9 \cdot 50\% = 5.45$ log points, and the aggregate business-stealing effect, since the average firm experiences business stealing from 50 percent of competitors, is $-8.6 \cdot 50\% = -4.3$ log points. This logic hinges on the model's prediction that both treated and untreated firms experience business-stealing effects. We present evidence on this prediction below.

TABLE 3—DIRECT AND INDIRECT EFFECTS: MAIN OUTCOMES

	log sales (1)	Profit (¥10,000) (2)	log num of employees (3)	log wage bill (4)	Fixed assets (¥10,000) (5)	log material cost (6)	Shutdown (7)
Post × Treated	0.099 (0.035)	12.64 (3.099)	0.075 (0.029)	0.101 (0.029)	5.468 (4.537)	0.077 (0.041)	−0.028 (0.010)
Post × Share Competitors Treated	−0.086 (0.041)	−9.478 (4.802)	−0.066 (0.038)	−0.069 (0.037)	−3.013 (4.558)	−0.050 (0.047)	0.001 (0.018)
Firm FE and Post	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of pure control	5.454	51.945	1.956	2.675	12.005	4.925	0
Observations	8,612	8,612	8,612	8,602	8,612	8,605	8,847
<i>Romano-Wolf p-values</i>							
Post × Treated	0.011	0.001	0.021	0.004	0.349	0.071	0.011
Post × Sh Comp Treated	0.057	0.064	0.093	0.071	0.607	0.404	0.912

Notes: Mean of pure control computed at baseline. Romano and Wolf p -values, which adjust for multiple hypothesis testing, are calculated with 1,000 bootstrap replications. Standard errors clustered at the market level.

treatment. Column 5 reports insignificant and small coefficients for fixed assets, suggesting that the loan was not used for traditional forms of capital. This result is consistent with the fact that most firms in our data are in the retail and service sectors and have low capital intensity. Column 6 reports effects on materials spending. The coefficients are comparable to those for employment, though less significant, and suggest that inputs were another margin of adjustment. Finally, column 7 shows that the treatment had a positive direct effect, and no indirect effect, on survival. The positive direct effect, though formally outside of our model, is consistent with the logic that borrowing improved firm performance. One reason for the lack of an indirect effect may be that business stealing is spread out over multiple competitors and thus has a smaller impact per firm. Overall, consistent with our model, Table 3 documents positive direct and negative indirect effects of the intervention.¹⁷

Linearity of Indirect Effect.—Our regression specification imposes the restriction that the business-stealing effect is a linear function of the share of competitors treated. As shown in Proposition 2, this restriction is an implication of our model. To evaluate its empirical validity, we estimate more flexible specifications that allow the indirect effect to vary with the intensity of exposure. We group firms in treated markets into three terciles based on the share of their competitors treated and replace the linear indirect effect in the estimating equation with indicators for these terciles (interacted with *Post*). Panels B–D of Figure 3 plot the estimates for three main outcomes: log sales, profit, and log employment. We also plot the straight line predicted by the linear specification of Table 3. The point estimates line up reasonably closely with the line in all three cases, as well as for take-up, shown in panel A. Given the standard errors, we cannot reject the linear specification. We conclude that the linearity restriction implied by the model is consistent with the data.

¹⁷To adjust for multiple hypothesis testing, Table 4 also reports Romano and Wolf p -values (Romano and Wolf 2005), which closely line up with the conventional measures of significance.

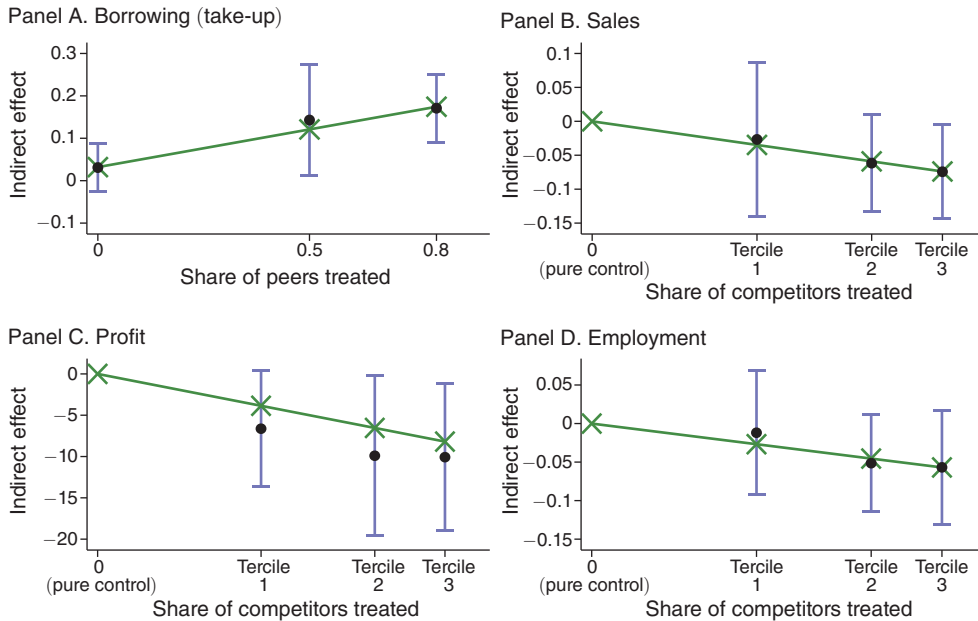


FIGURE 3. LINEARITY OF THE INDIRECT EFFECT

Notes: Panel A is based on column 3 of Table 2 and shows take-up intensity by treatment exposure. The green line is the prediction from the linear specification of column 2. Panels B, C, and D are based on heterogeneous effect regressions in which we group firms in treated markets into terciles based on the share of their competitors treated and use indicators for these terciles (interacted with Post) to estimate the indirect effect. The green line shows the prediction of the linear model from Table 3.

Specification Checks and Robustness.—We next show that our results are robust to some natural alternative empirical choices. First, we note that our second explanatory variable, the share of competitors treated—although, based on our discussion in Section IIC, it is as good as random—is not a treatment arm. It is natural to explore the alternative empirical approach in which we regress outcomes directly on the treatment arms. We present results from that specification for our three main outcomes in online Appendix Table A5. We find that, although power is somewhat lower—unsurprisingly, as we now have twice as many independent variables—the results closely correspond to those in Table 3. In particular, treated firms perform better than untreated firms in treated markets, while untreated firms in treated markets perform less well than untreated firms in control markets.

Second, although we use difference-in-differences as our main specification because that is what we derived from the model, in online Appendix Table A5, we show that our main results are similar when we estimate an ANCOVA specification (McKenzie 2012), in which we use the follow-up data and control for the baseline outcome.

A third issue is that we do not observe subsequent outcomes for firms that shut down. In our main approach, we include these firms in the regression only in those years in which they are in operation. We show in online Appendix A.1.6 that this approach is valid as long as our explanatory variables are orthogonal to the error term in the sample of firms that do not exit or attrit. We present evidence corroborating this

condition by showing in online Appendix Table A3 that our explanatory variables are orthogonal to baseline characteristics in the subset of firms that remain in the sample until 2016 or 2020. Nevertheless, we also pursue the alternative approach in which we keep firms that shut down in the sample and set their outcomes to zero after they shut down. We report the results from this approach, for sales, profit, and employment, in online Appendix Table A6. Because we can no longer take logs, we measure the outcomes in levels and with the inverse hyperbolic sine (IHS). Although power is marginally reduced, we find that neither the change in the sample nor the change in the way we transform the outcomes (level or IHS) affects our qualitative results.

Fourth, we note that, although such a force is absent from the model, in reality, firms could be stealing business from in-sample firms in other markets. Not accounting for such effects may lead to misspecification. However, such effects are unlikely in our context; as noted in Section IA, no county has two markets that are close competitors in terms of broad product category, and counties are far from each other, with a minimum distance between them of over 65 kilometers.

Heterogeneity by Treatment Status and over Time.—To gather further evidence on internal validity, we explore some model-implied heterogeneities in our main specification. First, we ask whether indirect effects are heterogeneous by the firm’s treatment status. Proposition 2 implies that, to a first-order approximation, business stealing should affect treated and untreated firms equally. To check this, we include in the regression the interaction of *SharePeersTreated* with both *Treated* and *Untreated*, resulting in a natural “saturated” specification. We report the results for our three main outcomes in online Appendix Table A7. The coefficients are insignificant, except for the negative indirect effect on untreated firms. The signs and magnitudes of the direct effect, and of the indirect effect on untreated firms, are similar to those in our main specification in Table 3. The magnitude of the indirect effect on treated firms is smaller than, but not significantly different from, that of the indirect effect on untreated firms. We conclude that we lack the power to separately identify indirect effects by treatment status and that we cannot reject our model-implied main specification.

Second, we ask how the direct and indirect effects vary between the midline and the endline. We do this by decomposing the interactions with *Post* into interactions with *Midline* and *Endline*. The results, reported in online Appendix Table A7, show that impacts at midline and endline are statistically not distinguishable from each other and that those at endline are more significant. These results suggest that, consistent with the model, the treatment had a nontransitory effect on business outcomes.

C. Intermediate Outcomes and Mechanisms

What mechanisms drove the direct and indirect effects? Our model identifies two mechanisms, as depicted in Figure 1: improvements in quality and variety and reductions in cost and price, both of which are expected to improve consumer experience and generate a reallocation of demand. To gather evidence on these mechanisms, we investigate impacts on intermediate outcomes.

TABLE 4—DIRECT AND INDIRECT EFFECTS: BUSINESS OUTCOMES

	log num of clients (1)	Renovation (2)	New product (3)	Quality of labor (4)	Supplier change (5)	Stocking period (unit: month) (6)	Inventory management (7)
Post × Treated	0.083 (0.032)	0.243 (0.020)	0.231 (0.018)	0.111 (0.031)	0.110 (0.026)	0.571 (0.086)	0.132 (0.021)
Post × Share Competitors Treated	−0.071 (0.034)	−0.049 (0.030)	−0.047 (0.019)	−0.022 (0.039)	0.031 (0.033)	0.026 (0.111)	0.014 (0.027)
Firm FE and Post	Yes	Yes	Yes	No	No	No	No
Mean of pure control	3.067	0.134	0.078	0.166	0.153	1.281	0.291
Observations	8,612	8,612	8,612	1,804	2,714	2,714	2,714
<i>Romano-Wolf p-values</i>							
Post × Treated	0.014	0.001	0.001	0.002	0.001	0.001	0.001
Post × Sh Comp Treated	0.067	0.209	0.015	0.742	0.742	0.742	0.742

Notes: Column 4 reports a cross-sectional regression using the 2020 follow-up data. Columns 5–7 report cross-sectional regressions using the endline data. In columns 4–7, we include controls for baseline employment and for the county and broad product category of the market. In columns 2, 3, and 5, dependent variable is an indicator. In column 4 dependent variable is the share of workers who completed high school; in column 7 it is an indicator for the firm having digitalized inventory records or a designated area for inventory storage. Mean of pure control is computed in the first wave in which the outcome was measured. Romano and Wolf *p*-values are calculated using 1,000 bootstrap replications. Standard errors clustered at the market level.

Business Outcomes and Practices.—We begin by looking at impacts on various business outcomes and practices, shown in Table 4. In column 1, we find a significant positive direct effect and a significant negative indirect effect on the firm’s number of clients. This result is consistent with the model’s logic of demand reallocation: as borrower firms improve, clients migrate from firms that have many borrower competitors to firms that are borrowers. The result is also consistent with the findings of Einav et al. (2021), who document the importance of customers in accounting for retail sales variation.

The rest of Table 4 seeks to identify dimensions of improvement that may have driven this reallocation. In columns 2–4, we study outcomes related to the mechanism of improvements in quality and variety. Column 2 reports impacts on renovation. We estimate a large and significant direct effect—the treatment increases the probability of renovation by 24 percentage points—and no indirect effect. Column 3 reports impacts on the introduction of new products. Here, too, we estimate a large and significant direct effect, again with almost a quarter of treated firms being impacted, and a much smaller indirect effect. Column 4 reports impacts on the quality of labor, measured as the share of workers in 2016 who had finished high school. We collected these data in our 2020 follow-up survey. Because we only have these data for one year, we estimate a cross-sectional regression.¹⁸ Again, we find a large positive direct effect and no indirect effect.

¹⁸In this and most subsequent cross-sectional regressions, we include controls for baseline firm employment and for the county and broad product category of the market. Note that the variable labels in Table 4 are correct because *Post* = 1 for the year 2016.

In columns 5–7, we explore outcomes related to the mechanism of reductions in marginal cost, which we collected in the 2016 endline survey. We focus on improvements in inventory and supplier management. Column 5 shows a positive direct effect on an indicator for the firm switching to a new supplier. One possible explanation is that the loan allowed the firm to place larger orders, which may have required changing suppliers. Consistent with this explanation, column 6 shows a positive direct effect on the stocking period, i.e., the number of months between episodes of restocking the store. This may have reduced marginal cost because large orders often come with a discount.

Column 7 shows an improvement in the quality of inventory management, measured with an indicator that equals one if the firm either had digitized inventory records or had a designated area for inventory storage.¹⁹ These effects could come about if the better-educated workers digitized records or the renovation created space for inventory. There are no detectable indirect effects in any of these cost-related outcomes.

In summary, the results in Table 4 provide evidence that, consistent with the model's mechanisms, the loan enabled investments to improve quality and reduce cost. The weak indirect effects suggest market-wide gains. These results are consistent with our survey of borrowers at endline, in which we asked them to describe what they primarily used the loan for and grouped the answers into categories. The three categories mentioned most frequently were renovation and increase in scale of operations (75 percent), purchase of inventory or inputs (50 percent), and starting new projects or introducing new products (34 percent), in line with our results on renovation, supplier/inventory management, and new product introduction.

Consumer Experience.—We next investigate whether—as the model predicts—the loan-induced changes in business practices affected consumer well-being. To do this, we look at impacts on different dimensions of consumer experience, including both price and quality experienced by consumers. We collected data on these outcomes in our short follow-up survey in 2020. We think that these data are an important part of our contribution.

To measure price, we asked firms to report the average price during 2016 of the product they sold that had the highest sales share. More precisely, in the 2016 survey, we had asked firms to list their top three main products ranked by their sales share. In 2020, we asked them for the average price in 2016 of the top product on that list. Our experience was that managers typically answered after consulting their records to get the exact historical price. When the manager was not available during our visit, we phoned to collect these data. When we were not able to reach the manager on the phone, we asked the market office, who had collected quarterly data on the price of the firm's main product as part of their regular data collection activities, and computed the average over 2016 of their quarterly records.²⁰ In this way, we were able to gather data on the 2016 price for 2,482 firms.

¹⁹Separately estimating the impacts on digitization and inventory storage yields similar results. Because answers to the two questions are positively correlated, we decided to group them into a single measure.

²⁰The market office collected these data in part to share with clients who might approach them.

We are able to validate these retrospective data because we had collected the price of the same main product in the 2016 survey for a subset of 262 firms. Thus, for this subset of firms, we have two observations about their 2016 price: one contemporaneous, the other retrospective. Regressing the retrospective 2016 log price (from the 2020 survey) on the contemporaneous 2016 log price (from the 2016 survey) yields a coefficient of 0.99 and an R^2 of 0.95, suggesting that retrospective price data are reliable.

To measure the quality experienced by consumers, we asked a random client physically present in the store to evaluate it along a number of dimensions. These included whether they received advice from the seller (yes/no) and their evaluation on a five-point scale of service quality, the shopping environment, value for money, and their overall satisfaction. We were able to collect these data for the 1,804 firms that we found open during our visit in 2020.

We validate these consumer satisfaction measures in two ways. First, we regress the first four evaluated dimensions on the fifth, i.e., the buyer's overall satisfaction. The coefficient in each specification is highly significant ($p < 0.001$), and the average R^2 is 0.27, suggesting that these measures contain relevant information about the firm. Second, to check whether the consumer's assessment of value for money contains information about price, we regress the log price in 2016, controlling for the main product of the firm, on the value-for-money score. Note that the price and value-for-money score are collected from different respondents. We obtain a significant negative coefficient ($p < 0.02$), implying that the consumer's value-for-money evaluation contains information about the price. Naturally, we expect that our measures also contain measurement error. Since we use them as outcomes, independent measurement error should not bias our estimates.

Table 5 reports the impacts on these outcomes. We estimate cross-sectional regressions because the data are only available in one wave. In column 1, we look at the price of the firm's main product. We find a significant negative direct effect of about 6 log points and an insignificant and small indirect effect. In columns 2–6, we look at the different quality and price-adjusted quality dimensions: advice from sellers, service quality, the shopping environment, value for money, and overall satisfaction. For all of them, we find a significant positive direct effect and a much smaller and less significant indirect effect. These results support our interpretation that the changes in business practices improved consumer experience.

In summary, Tables 4 and 5 suggest that—consistent with the model's mechanisms depicted in Figure 1—the loan enabled firms to improve quality and reduce cost, leading to an increase in consumer welfare and a reallocation of demand. Our findings also point to the importance of quality for retail firm performance, consistent with the results of Hottman, Redding, and Weinstein (2016) about the role of product appeal in the retail sector.

Tables 4 and 5 also highlight a pattern not implied by our model: a negative indirect effect on new product introduction and some measures of consumer experience. Indeed, our model only implies negative indirect effects for revenue and profit but not for business practices or consumer satisfaction. One explanation for the observed pattern may be that increased competition discouraged innovation, as predicted by the model of Aghion et al. (2005). Another explanation for consumer experience (but not new product introduction) may be that consumers in treated

TABLE 5—DIRECT AND INDIRECT EFFECTS: CONSUMER EXPERIENCE

	log price (1)	Advice from sellers (2)	Service quality (3)	Shopping environment (4)	Value for money (5)	Overall satisfaction (6)
Treated	−0.062 (0.026)	0.238 (0.035)	0.753 (0.0950)	0.991 (0.0969)	0.618 (0.077)	0.836 (0.060)
Share Competitors Treated	0.003 (0.040)	−0.098 (0.046)	−0.175 (0.120)	−0.345 (0.128)	−0.197 (0.086)	−0.231 (0.095)
Mean of pure control Observations	1.602 2,482	0.401 1,804	−0.004 1,804	0.043 1,804	0.024 1,804	0.021 1,804
<i>Romano-Wolf p-values</i>						
Treated	0.033	0.001	0.001	0.001	0.001	0.001
Sh Comp Treated	0.920	0.054	0.175	0.01	0.036	0.028

Notes: In column 1, sample is all firms we could reach at the 2020 follow-up to collect price data. In columns 2–6, where the outcome is based on a consumer’s evaluation, sample is all firms we found open during the 2020 follow-up. In column 2 outcome is an indicator; in columns 3–6 it is the z -score of the corresponding dimension of consumer evaluation. In all regressions, we include controls for baseline employment and for the county and broad product category of the market. Romano and Wolf p -values are calculated using 1,000 bootstrap replications. Standard errors clustered at the market level.

markets, having experienced some treated stores, had higher expectations and gave lower marks to untreated firms.²¹

Other Outcomes.—In online Appendix Table A8, we investigate impacts on other intermediate outcomes. First, we look at the use of trade credit with suppliers and with clients and find a positive direct and a negative indirect effect for both. One possible explanation is that trade credit tracks the intensity of business activities. Second, we look at advertising spending and find a positive direct effect but no indirect effect, indicating that advertising may have contributed to the reallocation of demand. However, the effect is quantitatively small (about ¥2,100 or US\$320), suggesting that the role of advertising was minor. Third, we find no direct or indirect effects on a markup measure computed as the ratio of revenue to cost, or on rental payments, suggesting that these channels did not contribute to our findings. We also find no impact on the log number of suppliers of the firm, suggesting that the switch to new suppliers documented above crowded out preexisting suppliers.

Alternative Mechanisms.—Because all our data come from self-reported surveys, a concern is that some of our results may be explained by experimenter demand. We have two pieces of evidence that address this concern. First, in the endline survey, we had our enumerators ask firms to directly show the value of sales in their books. We think that these book value data have about the same quality as administrative data and are unlikely to be affected by experimenter demand. Over 90 percent of firms surveyed at the endline showed us the book value, and nonresponses were

²¹It may be puzzling that we observe these negative effects on untreated firms even though some of them borrowed due to information diffusion. But untreated firms in treated markets borrowed on average 11 months later than treated firms, so that the effects of borrowing may not be fully visible here. Our IV strategy accounts for these forces.

balanced across the treatment arms. In online Appendix Table A9, we report impacts on the log of the book value of sales. We use a cross-sectional regression because the data are only available at the endline, and estimate effects that are less significant but numerically very close to those in Table 3. When we also include the baseline self-reported value of sales—collected before the treatment and therefore plausibly immune to experimenter demand—the estimates become significant. Moreover, when we use as the outcome the difference between the log of self-reported sales and the log of book sales, we find null effects. These results are direct evidence against experimenter demand. A second piece of evidence is the set of results on consumer satisfaction, which are based on the evaluations of customers who were not aware of the firms' treatment status. We conclude that experimenter demand is unlikely to drive our results.

A second alternative mechanism is that the regular visits by loan officers may have affected treated firms through a channel different from access to the loan. Two concrete channels seem salient: the visits may have increased managers' financial knowledge or improved their connection to the bank. To address the financial knowledge channel, we exploit data we collected in the endline survey about managers' financial knowledge using five questions described in the online Appendix (e.g., whether our partner bank offer loans to micro and small enterprises). Online Appendix Table A9 reports essentially null effects on an index of financial knowledge created by standardizing the average response to these questions. To address the connection channel, we use data we collected in the endline survey about the manager's assessment on a five-point scale of how difficult it is to get a loan. Online Appendix Table A9 reports essentially null effects on the standardized value of the response. We conclude that these channels probably did not contribute to the effect of the treatment.

D. Market-Level Effects

We turn to present regressions that assess the impact of the intervention at the market level. This analysis plays two roles. First, it validates our earlier arguments about how the firm-level estimates of direct and indirect effects aggregate to the market level. Second, it sets the stage for the model-based calculations of market-level welfare effects that follow in Section IV.

In Table 6, we report regressions measuring the impact of the share of firms treated on market-level outcomes. In these regressions, each observation is a market in a survey wave, and the outcome is an aggregate of firm-level outcomes across all firms in the market in that wave. Columns 1 to 3 show insignificant effects on the market-level revenue, profit, and employment. These results are consistent with the opposite-signed direct and indirect firm-level effects and reinforce the interpretation that the loan program led to within-market reallocation but no detectable gains in market-level producer surplus.

The remaining columns report impacts on a set of outcomes for which the firm-level regressions showed a direct effect substantially larger than the indirect effect: the shutdown rate, the renovation rate, the product introduction rate, labor quality, price, and customer satisfaction, the latter measured as a *z*-score at the market level. The precision of the effect for price is just below the conventional level of

TABLE 6—MARKET-LEVEL OUTCOMES

	log market revenue (1)	Market profits (2)	log market employment (3)	Shutdown rate (4)	Renovation rate (5)	Product intro rate (6)	Quality of labor (7)	log price (8)	Customer satisfaction (9)
Post × Share Market Treated	0.065 (0.047)	55.72 (127.98)	0.052 (0.039)	−0.072 (0.027)	0.162 (0.030)	0.146 (0.020)	0.052 (0.031)	−0.066 (0.045)	0.385 (0.061)
Market FE and Post	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
Mean of pure control	9.12	2,089.54	5.56	0	0.11	0.06	0.2	1.6	0.025
Observations	234	234	234	234	234	234	78	78	78

Notes: Quality of labor is the share of employees in the market with at least high school education. Log price is revenue-weighted average of the firm-level log price. Customer satisfaction is the z -score of the market-level average of overall customer satisfaction. Mean of pure control computed in the first wave in which outcome is measured. Standard errors clustered at the market level.

significance ($p = 0.13$), and the effect for labor quality is insignificant. However, the effect for every other outcome is significant.²² These results are consistent with the interpretation that—despite the insignificant effects on producer surplus—the treatment generated market-wide gains in price-adjusted quality and that consumers valued these gains.

Our model can rationalize the insignificant market-level revenue and the significant market-level quality effects with an elasticity of substitution between markets, θ , which is close to one. With such a θ , equation (2) implies that preferences over markets are approximately Cobb-Douglas, so that the relative expenditure share of each market is essentially constant. Intuitively, in this case demand at the market level is sufficiently inelastic that improvements in price or quality generate additional demand only to the extent that keeps market-level expenditure constant, but there is no reallocation from outside markets.

Our takeaway from these results is not that market-level growth is zero but that market-level growth can be meaningfully lower than the firm-level treatment effect because some of that treatment effect reflects reallocation within the market. An intuitive possibility is that the demand elasticity θ , and market-level revenue along with it, will increase once enough consumers outside the market learn about the improvements. Consistent with this mechanism, the market-level revenue effect, 6.5 log points, is not small. But it is insignificant and is thus also consistent with a market-level growth of zero. We are not able to study this issue in more depth with our data.

E. Geography and Competition

So far, we have assumed that all peers in the market induce information diffusion of the same intensity, and all competitors in the market induce business stealing of

²²Labor quality in the table is an employment-weighted average that measures the labor quality of the average worker. Using the unweighted average instead, which measures the labor quality of the average firm, would give significant results. In addition to these reported results, the market-level impacts on all other measures of consumer experience from Table 5 are positive and significant.

the same intensity.²³ However, it is possible that these indirect effects vary with geographic distance and vary differently for competitors and noncompetitors. To explore these effects, we use the fact that our data have information on the four closest neighbors of each firm. We call those four neighbors “local” to the firm and create four categories of peer firms in the market: local competitors, local noncompetitors, nonlocal competitors, and nonlocal noncompetitors. For example, local competitors are those firms among the four closest neighbors that have the same specialized product category, and nonlocal noncompetitors are those firms in the market that are not among the four closest neighbors and have a different specialized product category.

We estimate regressions in which the key right-hand side variables are the share of treated firms within each of these four peer groups.²⁴ One issue with this approach is that, when the size of a peer group is zero, the share is undefined. This is the case for 42 percent of firms for the local noncompetitors group, 11 percent of firms for the nonlocal competitors group, and less than 1 percent of firms for the other two groups. We “dummy out” these cases by including indicators of the firm having zero peers in each group (interacted with *Post*) and setting the share treated to zero in these cases.²⁵ Thus, our regression measures the impact of the share of peers treated for the subsample of firms for which the peer group is nonempty.

Whenever a peer group is nonempty, the randomness of the treatment assignment ensures—as with our main estimating equation—that the share treated is as good as randomly assigned. Indeed, balance tests (not reported) confirm that the share treated in each peer group is uncorrelated with the firm’s baseline characteristics. Thus, our regressions measure the causal effect of the share of firms treated in each group. However, because location and competition status are not randomly assigned, differences between the coefficients across peer groups are not necessarily driven by the causal effect of distance and competition status.²⁶ For this reason, our interpretation of the differences between the coefficients should be viewed as suggestive.

We first investigate heterogeneity by geography and competition status in the information diffusion effect. Table 7 reports regressions of loan take-up on the share of peers treated in each of the above-defined four groups. Column 1 focuses on the sample of treated firms and shows that for them, the likelihood of taking out the loan does not vary with the share treated in any of the groups. This seems plausible; treated firms do not need to rely on information diffusion because they know about the loan directly from the loan officer.

Column 2 reports analogous estimates for the sample of untreated firms. Here, the results are more interesting. Beginning with local peers, we find a small and

²³The latter statement is conditional on size; our approach incorporates the idea that larger firms generate more business stealing.

²⁴We are able to construct these variables even though our sample contains only half of the firms in the market, and hence does not cover about half of the listed neighbors, because we know the treatment status of all firms.

²⁵Omitting these cases instead leads to slightly noisier but qualitatively similar results.

²⁶For example, if managers locate businesses close to their friends, then the effect of local competitors getting the treatment may be driven by friendship. That said, firm location is decided by the market office based on factors such as space availability. In a subsample of 140 firms for which we know the managers’ 5 best friends in the market, we find that the share of local competitors, local noncompetitors, nonlocal competitors, and nonlocal noncompetitors who are friends is 1 percent, 1 percent, 3 percent, and 7 percent. Thus, friends grouping together may not be a major concern.

TABLE 7—EFFECTS ON BORROWING BY PEERS' LOCATION AND COMPETITION STATUS

Dependent variable: Sample:	Borrow with new loan product	
	Treated (1)	Untreated (2)
Share Local Competitors Treated	−0.025 (0.039)	−0.019 (0.044)
Share Local Noncompetitors Treated	0.058 (0.065)	0.110 (0.050)
Share Nonlocal Competitors Treated	0.004 (0.092)	0.098 (0.058)
Share Nonlocal Noncompetitors Treated	−0.047 (0.147)	0.069 (0.078)
Mean of pure control	0.031	0.031
Observations	1,256	1,525

Notes: Sample is all firms in the 2016 endline survey, which are the firms for which we have the neighbor data. Column 1 includes treated, column 2 untreated firms. Indicators for no peers in the relevant group, interacted with Post, are included in both specifications. Standard errors clustered at the market level.

insignificant effect of the share of local *competitors* treated but a large and significant effect of the share of local *noncompetitors* treated. A possible mechanism explaining this difference—shown to be active in different contexts by Cai and Szeidl (2018) and Hardy and McCasland (2021)—is that firms prefer to withhold business-relevant information from their direct competitors. Indeed, our results on business stealing confirm that withholding information about the loan from competitors is in the interest of the firm. In contrast, for local noncompetitors, there is no risk of business stealing, and improvements made by treated neighbors may attract more shoppers to the neighborhood.

Turning to nonlocal peers, we find a significant effect of the share of nonlocal *competitors* treated. This is surprising, as there could be business stealing by these firms as well. A possible explanation is that, because they are located farther away, these peers do not directly compete with the firm; however, because they are in the same business, they do share information. Consistent with this logic, we show below that business stealing by these competitors is weaker. Finally, we find a noisily estimated positive effect of the share of nonlocal *noncompetitors* treated. This is the largest of the four peer groups, so the noisy estimate may indicate either that there is no diffusion from these firms or that there is diffusion from a subset, e.g., from those who are friends with the manager of the firm.

Taken together, these heterogeneous effects support our interpretation that the borrowing spillover reflects information diffusion and highlight the role of agents' incentives to talk in shaping social learning and technology adoption (Banerjee et al. 2018; Chandrasekhar, Golub, and Yang 2018).

We next investigate heterogeneity by geography and competition in the indirect effect on firm performance. Panel A of Table 8 reports estimates of the impact on three main outcomes of the share treated in each of the four groups. We begin with peers who are competitors. Consistent with the logic that the main competitors of a firm are local, the share of treated among *local* competitors has a significant negative effect on log sales and profits. By contrast, the share of treated among *nonlocal*

TABLE 8—EFFECTS ON MAIN OUTCOMES BY PEERS' LOCATION AND COMPETITION STATUS

	Panel A. All firms			Panel B. Treated and pure control		
	log sales (1)	Profit (¥10,000) (2)	log num of employees (3)	log sales (4)	Profit (¥10,000) (5)	log num of employees (6)
Post × Treated	0.095 (0.041)	11.91 (2.753)	0.078 (0.031)	0.048 (0.142)	-6.178 (11.83)	0.030 (0.058)
Post × Share Local Competitors Treated	-0.122 (0.054)	-11.84 (4.783)	-0.064 (0.039)	-0.029 (0.068)	-4.295 (4.672)	0.018 (0.038)
Post × Share Local Noncompetitors Treated	0.190 (0.049)	14.70 (4.981)	0.062 (0.030)	0.169 (0.056)	17.79 (5.744)	0.015 (0.027)
Post × Share Nonlocal Competitors Treated	-0.113 (0.072)	-10.31 (11.62)	-0.038 (0.049)	-0.171 (0.126)	-1.600 (13.36)	0.008 (0.062)
Post × Share Nonlocal Noncompetitors Treated	0.052 (0.062)	7.569 (14.68)	0.002 (0.050)	0.085 (0.094)	12.42 (15.07)	-0.043 (0.047)
Firm FE and Post	Yes	Yes	Yes	Yes	Yes	Yes
Mean of pure control	5.45	51.95	1.96	5.45	51.36	1.96
Observations	8,220	8,220	8,220	6,967	6,967	6,967

Notes: Panel A uses firms in the 2016 endline, for which we have the neighbor data. Panel B uses the subsample that are treated or in the pure control group. Indicators for no peers in the relevant group, interacted with Post, are included in all specifications. Mean of pure control computed at baseline. Standard errors clustered at the market level.

competitors has a smaller and imprecisely estimated effect on all three outcomes. This result helps rationalize why we observe information diffusion from nonlocal, but not from local, competitors.

The most interesting result of Table 8 concerns noncompetitors. We find, unexpectedly, that the share of local noncompetitors treated has a *positive* effect on firm performance. Being positive, and coming from noncompetitors, this effect cannot be driven by business stealing. Two other explanations seem possible. First, the effect could be driven by the diffusion of information: treated noncompetitors may induce that the firm borrows and thus grows. Second, the effect could be driven by a new mechanism, the diffusion of demand: treated noncompetitors may attract more consumers to the neighborhood, who then also shop at the firm. Panel B of Table 8 attempts to distinguish between these explanations by estimating the same regression for the subsample of firms that are either treated or in pure control markets. For these firms, we do not expect the first channel to be active. Indeed, as we have seen in Table 7, treated firms do not experience information diffusion, while pure control firms—because they have no treated peers—cannot experience information diffusion. Table 8 shows that the positive indirect effect is preserved in this subsample, providing evidence in favor of the demand diffusion explanation.²⁷

In online Appendix A.1.7, we extend our model to incorporate the new mechanism of demand diffusion. This model generates the positive indirect effect on treated noncompetitors that we have documented and yields two new predictions, on

²⁷This is not to say that the information diffusion effect is inactive. We believe that it is active but weak because firms that borrow due to diffusion borrow about 11 months later, and hence, the impact of their loans was realized only after the midline survey.

which we present evidence in online Appendix Table A10. First, the share of treated noncompetitors should not have a positive effect on consumer satisfaction because the increased demand is due to “shopping around” rather than to improvements in quality and price. Our results confirm this. Second, demand diffusion itself should generate business stealing because the demand that results from shopping around must be reallocated from somewhere else. To test for this, we present regressions showing that firms with a higher share of competitors exposed to demand diffusion exhibit lower performance. Since these regressions also include the share of local competitors treated, which is correlated with the share of local competitors exposed to demand diffusion, we interpret the results as suggestive. But we still find them interesting because they imply a second-degree indirect effect—from the treatment to demand diffusion to business stealing—suggesting that indirect effects may accumulate over firm networks.²⁸

These results provide new experimental evidence on a demand externality that may be an important driver of the spatial concentration of retail establishments commonly observed around the world (Marshall 1920; Fujita and Thisse 1996; Leonardi and Moretti 2022). The demand externality we document acts between noncompetitors. However, because our markets are fairly specialized, even noncompetitors tend to be in the same broadly defined trade, suggesting that this externality can help explain the clustering of similar establishments.

IV. Combining Indirect Effects: Estimation and Welfare Evaluation

We now combine the direct and indirect effects of the loan program. We first present IV estimates that measure the impacts of the loan, accounting for both information diffusion and business stealing. We then combine the results with the model to evaluate the welfare impact of the program on both firms and consumers. Finally, we discuss the plausibility of some of our assumptions.

A. IV Estimates

We begin by measuring the impact of borrowing, as opposed to the treatment, on firm performance. Because the treatment effects on total borrowing are similar to the treatment effects on borrowing through the new loan (Table 2), we focus on evaluating the impact of borrowing through the new loan. For simplicity, in this analysis we ignore the heterogeneity in indirect effects and the new indirect effect of demand diffusion that we documented in Section III E.

Under these assumptions, Proposition 2 implies that we can estimate the direct and indirect effects of the loan using the IV regression (9), in which we instrument borrowing and the share of competitors that borrow with the treatment and the share of competitors that are treated. This IV accounts for the two indirect effects (information diffusion and business stealing) at different stages. Information diffusion is accounted for at the first stage, where the firm’s borrowing status can depend on the share of its competitors that are treated. Business stealing is

²⁸The accumulation of indirect effects, albeit of a different kind, plays an important role in theories of input-output networks, for example, in Acemoglu et al. (2012).

accounted for at the second stage, where the firm's performance can depend on the share of its competitors that borrow. This IV strategy assumes that the treatment only induces information diffusion about the loan, not about other practices. However, nontreatment induced information diffusion about other practices is allowed and is one reason we cluster standard errors at the market level.

To estimate the IV, we need to define which firm constitutes a borrower in which period. Because using the loan plausibly takes time, we classify untreated borrower firms—which borrow on average 11 months later than treated borrowers—as borrowers only at the endline ($s = \textit{Endline}$ in the model), effectively assuming that for these firms, the impact of the loan is only realized after the midline.²⁹

Table 9 reports the results. Columns 1 and 2 show the first stage for both explanatory variables. As expected, both instruments create variation in borrowing, while only the share of competitors that are treated creates meaningful variation in the share of competitors that borrow. The F -statistics for the first stages are over 50, suggesting that weak instruments are not a problem. Columns 3–5 show the second stage for three main outcomes. As in the reduced-form regressions, we estimate significant positive direct effects and significant negative indirect effects that are of comparable magnitude. The coefficients are larger than in the reduced-form regressions since here we evaluate the impact of borrowing, not of the treatment. The estimates imply that borrowing would increase sales by 32 log points and that increasing the share of competitors that borrow from 0 to 100 percent would reduce sales by 29 log points. We obtain similarly large direct and indirect effects for profit and employment.

We conclude that the qualitative findings from the IV are similar to those from the reduced-form regressions, validating our approach of using the reduced-form regression in most of our analysis. We now turn to use the IV estimates for welfare evaluation.

B. Welfare Evaluation: Strategy

We define the welfare gain from the loan program to be the total improvement in the welfare of firms and consumers in the markets, resulting from the direct, diffusion, and business-stealing effects, net of the interest cost of the loan. This measure of the welfare gain would approximate the societal welfare impact of the program if (i) there are no other indirect effects and (ii) the interest rate is a good measure of the social cost of capital. We discuss both assumptions in Section IVD; for now, we note that, even if they fail, our definition captures an important component of the societal welfare effect of the program.

We compute the welfare gain using the model of Section IIB, which omits firm exit and demand diffusion. We discuss below extensions that allow these channels and show that they have small effects on our results. In our model, the change in welfare from the loan program comes from a change in the consumer surplus and a change in the producer surplus. Let \tilde{R}_m^l denote the total revenue of firms in market

²⁹Classifying them as borrowers at midline generates slightly larger direct and indirect effects.

TABLE 9—EFFECTS OF BORROWING ON MAIN OUTCOMES: IV ESTIMATION

	First stage		IV		
	Borrow (1 = Yes, 0 = No) (1)	Share competitors borrow (2)	log sales (3)	Profit (4)	log number of employees (5)
Post × Treated	0.322 (0.034)	0.011 (0.006)			
Post × Share Competitors Treated	0.092 (0.023)	0.399 (0.038)			
Borrow			0.318 (0.127)	40.41 (9.698)	0.239 (0.07)
Share Competitors Borrow			−0.288 (0.134)	−33.09 (12.978)	−0.22 (0.082)
Firm FE and Post	Yes	Yes	Yes	Yes	Yes
F-statistics	50.10	58.84			
Observations	8,612	8,612	8,612	8,612	8,612

Notes: Borrow is an indicator for borrowing through the new loan product. For treated borrowers, it equals one at midline and endline; for untreated borrowers, who borrowed later, it equals one at endline. Mean of pure control computed at baseline. Standard errors are clustered at the market level.

m in period t absent the treatment. Then, the following result characterizes the impact of the loan program on the consumer surplus.

PROPOSITION 3: *To a first-order approximation, in period $t \geq s$ the impact of a loan program on consumer surplus is*

$$(10) \quad \Delta CS^t \approx \int_m \gamma [S_m \mu_m^T + \phi S_m (1 - S_m) \mu_m^D] \lambda_m^t \tilde{R}_m^t dm.$$

To understand equation (10), recall that the intervention reduces the price and increases the quality of borrower firms. The equation expresses consumers’ savings from purchasing the same (quality-adjusted) bundle that they would have purchased absent the intervention, but at the new (quality-adjusted) prices. By the envelope theorem, these savings—which could be spent on any good including the numeraire—are to a first-order approximation equal to the gain in the consumer surplus. The formula expresses these savings as the reduction in spending that results from borrower firms reducing their quality-adjusted prices. In particular, given take-up and diffusion, the share of firms that borrow equals $S_m \mu_m^T + \phi S_m (1 - S_m) \mu_m^D$, and these firms experience a reduction in their quality-adjusted price of γ . This reduction is relative to total market revenue \tilde{R}_m^t and is amplified because borrower firms are on average λ_m^t times larger than the average firm in the market.

To evaluate welfare impacts using this result, we make two simplifying assumptions. First, we focus on a counterfactual intervention in which a constant share S of firms are treated in markets M . Second, we set aside cross-market heterogeneity and evaluate impacts for the average market. In particular, we approximate the market-level take-up and diffusion intensities μ_m^T and $\phi \mu_m^D$ with their sample averages μ^T and $\phi \mu^D$ estimated in column 2 of Table 2. We measure market revenue absent the treatment, \tilde{R}_m^t , by multiplying average firm revenue at midline and endline in pure control markets, \bar{R}_C , with the number of firms in the market n_m . And,

because at baseline the average revenue of borrowers relative to all firms is about 1.1 and not significantly larger than 1, we proxy the market-level selection intensity λ_m^t with 1 in all markets.

To estimate the deep parameter γ , note that by Proposition 2, the regression coefficient measuring the direct effect of the loan on revenue (in column 3 of Table 9) is an estimator for $\gamma(\sigma - 1)$. We infer γ from this coefficient by calibrating σ based on the literature. Note that the higher the σ , the lower the implied γ , and the lower the estimated welfare gains. Recent work estimating the retail elasticity of substitution includes Atkin, Faber, and Gonzalez-Navarro (2018), who find σ in the range of 2.3–4.4, and Dolfen et al. (2019), who find σ in the range of 4.3–6.1. For our main results, we set a conservative $\sigma = 6$, which is at the high end of these ranges and close to the inverse of the average profit-to-sales ratio in our baseline sample, 5.9, which would be equal to σ in our model. For robustness, we also consider $\sigma = 4$ and $\sigma = 8$ in online Appendix A.2.

Welfare Formulas.—We are ready to state our formulas for the welfare impact of the program. Given our assumptions, Proposition 3 implies that the impact of the program on consumer surplus, normalized by the total mass of firms in treated markets, is

$$(11) \quad \frac{\Delta CS}{\sum_M n_m} \approx \frac{\zeta_R}{\sigma - 1} [S\mu^T + \phi\mu^D S(1 - S)] \cdot \bar{R}_C,$$

where ζ_R is the IV estimate of the direct effect of borrowing on log revenue from Table 9. The impact on producer surplus, that is, profits, can be estimated directly using the coefficients of the profit IV regression:

$$(12) \quad \frac{\Delta PS}{\sum_M n_m} = (\zeta_{\Pi} + \xi_{\Pi}) [S\mu^T + S(1 - S)\phi\mu^D],$$

where ζ_{Π} and ξ_{Π} are the direct and indirect profit effects from Table 9. Intuitively, each firm gets the loan with probability $S\mu^T + S(1 - S)\phi\mu^D$, and if it does, experiences an average profit gain of ζ_{Π} . Further, each firm experiences a business-stealing effect from the share of borrowing competitors $S\mu^T + S(1 - S)\phi\mu^D$ of magnitude ξ_{Π} . As above, μ^T and $\phi\mu^D$ come from the take-up regression in column 2 of Table 2. These consumer and producer surplus effects are our main measures of the welfare gain from the loan.

Return on Capital.—To make comparisons with other credit interventions, it is helpful to compute the return on capital. Following Banerjee and Duflo (2014), we define the private return to capital as the return that business owner would earn from injecting capital into the business. Equivalently, this definition measures the return accumulating to the bank and the borrower firm as a result of the loan. We can analogously define the social return on capital, which in addition accounts for the effect of the loan on market competitors and consumers.

A simple calculation for the private return on capital is to take the estimated effect of the loan on profit from Table 9 of ¥404,100 and divide it by the average loan

amount of ¥290,000. This gives a net private return of 139 percent. However, this calculation uses profits measured starting two years after the intervention—midline is two years after baseline—and it is plausible that in the first year the profit gains were lower. We therefore develop an approach to compute returns while setting the gains in the first year to zero.

Our approach assumes that all firms in a market are treated, so that there are no diffusion effects. We then proceed in two steps, which are explained in detail in online Appendix A.1.5. First, we measure the average *yields* of the loan as estimated in our data. This is essentially the calculation we conducted in the previous paragraph, with two modifications: we adjust for the interest rate and the default rate that affect the bank's earnings, and we also compute yields that measure the additional contributions of business stealing and the consumer surplus. Second, we make assumptions about the dynamics of yields: that all yields in the first year are zero, that the yield in year 2 is what we computed, and that in subsequent years all yields depreciate at an annual rate $d = 0.10$.³⁰ We then compute the internal rate of return associated with this time path of yields. We use this procedure to compute the private and the social return, as well as to decompose the social return into the contributions of the private return, business stealing, and the consumer surplus.

We obtain confidence intervals for our welfare measures—including the producer surplus, the consumer surplus, and the return on capital—by bootstrapping the entire procedure, starting with the IV estimation, 1,000 times, drawing markets with replacement. As a result, the confidence intervals of our welfare estimates reflect the uncertainty about the IV coefficient estimates.³¹

C. Results

We first present the results on the welfare gain and then on the return to capital. Table 10 reports the implied impacts on the consumer surplus, the producer surplus, and total welfare. Bootstrapped standard errors and confidence intervals are reported in parentheses. All welfare gains are reported relative to the number of firms in the market (not relative to the number of firms treated) and are thus comparable across different treatment intensities S . The first two columns focus on the impact of treating all firms in the market. Column 1 reports impacts scaled by the profit of the average firm, while column 2 reports impacts in US dollars. The first row shows that treating all firms generates an insignificant gain in producer surplus amounting to 4 percent of profits. The second row shows that treating all firms generates a significant gain in consumer surplus amounting to 13 percent of profits, or about US\$11,000 per firm in the market. The total welfare gain is about 17 percent of profits, or close to US\$15,000 per firm. Because all firms are treated, we have no information diffusion effects.

The next two columns report the results when 50 percent of firms are treated. Here, we decompose both the consumer surplus and the producer surplus into a term

³⁰ Online Appendix Table A7 shows insignificant differences between effects at midline and endline, suggesting that a 10 percent depreciation is a conservative choice.

³¹ Because the internal rate of return is not defined for negative yields, for the approximately 10 percent of bootstrap draws in which the yield is negative, we set the return equal to the yield. This is a conservative choice for the confidence intervals of the implied returns.

TABLE 10—WELFARE GAIN ESTIMATES

Welfare gain per firm in market:	Treat all firms		Treat 50% of firms	
	Share of profit (%)	US\$	Share of profit (%)	US\$
Producer Surplus	4.1 (4.4) [-5, 12]	3,566 (3,904) [-4,174, 10,919]	2.0 (2.2) [-2, 6]	1,778 (1,952) [-2,210, 5,310]
Consumer Surplus	12.7 (4.6) [4, 22]	11,139 (4,022) [3,927, 19,618]	6.3 (2.3) [2, 11]	5,565 (2,011) [1,964, 9,809]
Spillover			2.4 (1.3) [0, 6]	2,087 (1,144) [317, 4,913]
Total	16.7 (7.3) [3, 32]	14,696 (6,415) [2,721, 28,027]	10.7 (4.9) [2, 21]	9,430 (4,280) [1,506, 18,297]

Notes: Bootstrapped standard errors in round brackets, and bootstrapped bias-corrected percentile confidence intervals in square brackets, are computed by bootstrapping our estimation procedure 1,000 times, drawing markets with replacement. US dollar values calculated using the average annual exchange rate during midline and endline (6.465).

measuring the impact of the loan program absent information diffusion and another term measuring the additional impact of information diffusion.³² We label the combined impact on the consumer and producer surplus of information diffusion as the spillover effect. As the table shows, the per firm effects on producer and consumer surplus are halved relative to the case when all firms are treated; but now information diffusion generates additional gains, which amount on average to 2.4 percent of firm profits or about US\$2,100 per firm in the market. These sizable gains raise the question of whether it may be optimal to treat only a subset of firms and leverage diffusion. In our model, the additional gain in consumer surplus from treating one more firm, if it is already exposed to diffusion from all its peers, is $\gamma(\mu^T - \phi\mu^D)\bar{R}_C$, about US\$4,000. This is almost certainly larger than the marginal cost of treatment and suggests that in our setting it is optimal to treat all firms.

These results have two main implications. First, they show that the welfare gain from the loan program was substantial and mainly driven by the consumer surplus. The fact that the incidence of the welfare gain is on consumers, not producers, suggests that policies introduced to improve industrial performance, even if they improve the affected businesses, may not achieve their goal but may nevertheless generate sizable welfare gains. The second implication is that accounting for, potentially multiple, indirect effects can be essential for the welfare evaluation of firm policies. In our setting, accounting for the direct and indirect effects on firms, while ignoring the effect on consumers, would imply that the program generated insignificant and small welfare gains, whereas also accounting for the effect on consumers implies large and significant welfare gains.³³

³²The former is obtained by setting ϕ to zero in (11) and (12), and the latter is obtained as the residual.

³³Online Appendix Table A11 replicates the welfare results for $\sigma = 4$ and $\sigma = 8$ and finds that our qualitative conclusions are robust: the program is estimated to generate large welfare gains mostly driven by the consumer surplus.

TABLE 11—RETURN TO CAPITAL DECOMPOSITION

Private Return (percent)	74.2 (12.9) [46, 98]
Business Stealing (pp)	-56.3 (23.4) [-104, -13]
Consumer Surplus (pp)	41.9 (13.6) [16, 70]
Social Return (percent)	59.8 (21.8) [11, 98]

Notes: Bootstrapped standard errors in round brackets, and bootstrapped bias-corrected percentile confidence intervals in square brackets, are computed by bootstrapping our estimation procedure 1,000 times, drawing markets with replacement. In draws with negative raw yields (about 10 percent of cases), we approximate the internal rate of return with the yield.

We now turn to the return on capital. Table 11 reports the implied returns and the decomposition. The private return is about 74 percent. Most of the private return is canceled by business stealing. The social return is about 60 percent and mostly driven by gains in consumer surplus. The large gap between the private return to capital and bank deposit rates, which were below the loan interest rate, suggests that some friction limits lending below the privately efficient level. And the fact that the social return of the loan is also very high suggests that due to this friction, large potential welfare gains are not realized.

It is useful to compare our results to estimates of the (private) return to capital obtained in other contexts. De Mel, McKenzie, and Woodruff (2008) estimate returns of 55–63 percent for microenterprises in Sri Lanka, while Banerjee and Duflo (2014) estimate a return of 105 percent for larger firms in India. Our private return of 74 percent for SMEs falls between these estimates.³⁴

Our analysis may contribute to understanding why the return to capital in developing countries is high. In our model, the private yield to borrower firms, of which the private return is a function, is proportional to the loan-induced improvement in quality-adjusted productivity (γ) times the potential for business stealing ($\sigma - 1$). Thus, for σ large enough, business stealing generates an amplification, through which even moderate improvements in productivity can result in high private returns. The logic of amplification also implies that small differences in the productivity-improving effects of the loan (γ) can generate large differences in the observed return to capital (through $\gamma \cdot (\sigma - 1)$), a possible mechanism predicting a large difference in returns between developing and developed countries.

³⁴Other experimental work in development estimates similarly high (uncompounded) annual returns to capital in microenterprises: McKenzie and Woodruff (2008) estimate 120–396 percent in Mexico, Fafchamps et al. (2014) estimate 180 percent in Ghana, and Field et al. (2013) estimate 156 percent in India.

D. Discussion of Assumptions for Welfare Evaluation

Other Indirect Effects.—Our welfare results assume that there are no other indirect effects. Omitted indirect effects may include effects on competitors from outside the market, effects on suppliers, and general equilibrium effects through income or the wage. On outside competitors, we expect at most weak effects: since there are no detectable impacts on market-level revenue, consumers do not seem to be reducing spending elsewhere. On suppliers, since we observe some switching, we do expect a reallocation effect. To the extent that treated firms are switching to higher-quality suppliers, this effect is likely to be welfare enhancing. As to business owners' income, because the impact on the producer surplus is small, we expect at most a small effect driven by differences in the propensity to spend. Finally, concerning the wage, because the direct and indirect effects on employment roughly cancel, we expect at most small effects. Thus, the omitted indirect effects on welfare are likely to be either approximately zero or marginally positive.

Interest Rate as Cost of Capital.—Our measure of the welfare gain net of interest approximates the true welfare gain under the assumption that the interest rate measures the cost of capital. This is a natural assumption for evaluating the private gain: since presumably the bank earns profits, the interest rate is a plausible upper bound for their private cost of capital. But for the social gain, the relevant measure is the social cost of capital, which may be different depending on the alternative use of capital. For example, if the bank, instead of lending to our firms, uses the capital to buy government securities, that itself could have indirect effects on other actors. Our measure does not account for this social opportunity cost, but comparing our social return to that in an alternative use of the capital would lead to valid conclusions about the relative desirability of these uses.

Alternative Loan Allocation.—A related question, motivated by Bertrand, Schoar, and Thesmar (2007) and Sraer and Thesmar (2020), is whether a different loan allocation could have generated larger welfare gains. This is not our main focus, but our expectation is that better allocations may exist. Heterogeneous effect regressions show that impacts were higher for firms that were larger or had a more educated manager (not reported), suggesting that targeting such firms could have increased efficiency.

Omitted Effects: Default, Exit, Demand Diffusion.—In the analysis we ignored the impact of loan default. Although we do not have direct data, the bank informed us that repayment rates for the two-year loan were over 98 percent, implying a default rate of less than 1 percent per year. Because the bank acts as a for-profit lender, given this low rate, we expect that interest payments covered the bank's cost of capital, including losses from default. Thus, computing the welfare gain net of interest payments accounts for default.³⁵

³⁵ Even if this is not the case, the quantitative effect of the 1 percent default rate on welfare is small.

The analysis ignored firm exit. In online Appendix A.1.6, we develop a model extension that incorporates exit and highlights two new effects. (i) Because some borrowers exit, the effects of the program on consumer and producer surplus are reduced over time. (ii) Borrowing increases survival, which, to the extent that borrowers are larger than average, increases the surplus. In the data the exit rate of borrowers is low, which implies that the first channel has little impact on our welfare results. In addition, at baseline borrowers are only slightly and insignificantly larger than the average firm, suggesting that the second effect is also plausibly small.

Finally, the analysis ignored demand diffusion. We note that including as controls the variables from online Appendix Table A10 measuring demand diffusion and its business-stealing effect has minor effects on the main coefficients (not reported). In addition, in online Appendix A.1.7 we develop a model extension that formalizes diffusion as a random reallocation of demand. Such a reallocation, to a first-order approximation, has no effect on welfare because the expected marginal utility of the goods that experience a demand increase is the same as that of the goods that experience a demand decrease.

V. Conclusion

We estimated the direct and indirect effects of a loan program. We found that borrower firms provided higher quality and variety and charged lower prices and that consumers valued these gains and reallocated their demand to borrowers. We estimated that these changes had a statistically undetectable effect on producer surplus but a large positive effect on consumer surplus. We also found indirect effects operating through the diffusion of information and the diffusion of demand. We now discuss some caveats and implications of these results.

We begin with external validity. Our primary interest is not in the external validity of the statistically zero effect on producer surplus but in that of the importance of business stealing, the underlying mechanisms, and the welfare implications. In this regard, a natural concern is that business-stealing effects may be especially strong in our spatially concentrated retail markets. It is certainly possible that business-stealing effects are different in other sectors, e.g., manufacturing. However, the retail sector appears to be spatially concentrated in both developing and advanced countries (Jensen 2007; Hardy and McCasland 2021; Leonardi and Moretti 2022). Thus, there is an important set of contexts where we may expect similar business-stealing effects.

We next compare our results in more detail to some key papers discussed in the introduction. Rotemberg (2019), in a credible but nonrandomized evaluation of a subsidy policy, finds, as we do, a large business-stealing effect. However, in contrast to our results on improvements in quality and cost, he does not find evidence on improvements in firm productivity. Possible explanations for this difference may be that he studies a subsidy rather than a loan program or that his data do not allow measuring impacts on product quality. Another difference is that Rotemberg (2019) does not explore impacts on consumers. However, the reallocation he documents suggests that consumers do benefit, so that there may be an impact on consumer surplus in his context too. McKenzie and Puerto (2021), in a randomized evaluation of a training program, do not find evidence on business stealing. One key difference

between their context and ours is that their business owners spent some of their income in the local market. This effect is likely absent in our specialized markets. Because it predicts that the treatment-induced gains may be spent on untreated peers, it can help close the gap between their results and ours. Finally, Drexler, Fischer, and Schoar (2014) and Calderon, Cunha, and De Giorgi (2020) find suggestive evidence consistent with business-stealing effects. We conclude that comparisons with prior work are consistent with the potential broad importance of business-stealing effects. Our contribution to this work is the randomized evidence on business stealing; the underlying mechanisms on quality, price, and consumer satisfaction; and the welfare implications.

Our results have implications for firm-level impact evaluations and for industrial policy. Concerning impact evaluations, our results suggest that a positive direct effect on firms, especially if these firms sell directly to consumers, likely reflects a gain in consumer surplus: after all, there must be a reason consumers increase their purchases. But the impact on producer surplus is less clear, as the direct effect may be partially offset by losses at untreated firms. Concerning industrial policy, our results suggest that policies introduced to improve industry outcomes such as employment, even if they increase employment at treated firms, may not achieve their goal due to business-stealing effects. However, these policies may increase consumer surplus and welfare.

Finally, our results provide new evidence that innovations that improve aggregate welfare can create economic losers. This force may be important for understanding the development process. In particular, our analysis suggests that, because in our model business stealing impacts all firms, including treated nonborrowers, the loan program had a negative effect on the *majority* of firms. If firms could have voted on the program while fully anticipating its implications, they should have voted it down. Of course, this is purely hypothetical, as such a collective action was not possible in our context. But our evidence does suggest that economic losers are an important consequence of development, and there may be contexts in which they act as a barrier.

REFERENCES

- Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi.** 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica* 80 (5): 1977–2016.
- Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt.** 2005. "Competition and Innovation: An Inverted-U Relationship." *Quarterly Journal of Economics* 120 (2): 701–28.
- Angelucci, Manuela, and Giacomo De Giorgi.** 2009. "Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?" *American Economic Review* 99 (1): 486–508.
- Angelucci, Manuela, Dean Karlan, and Jonathan Zinman.** 2015. "Microcredit Impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco." *American Economic Journal: Applied Economics* 7 (1): 151–82.
- Aronow, Peter M., and Cyrus Samii.** 2017. "Estimating Average Causal Effects under General Interference, with Application to a Social Network Experiment." *Annals of Applied Statistics* 11 (4): 1912–47.
- Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro.** 2018. "Retail Globalization and Household Welfare: Evidence from Mexico." *Journal of Political Economy* 126 (1): 1–73.
- Atanasio, Orazio, Britta Augsburg, Ralph De Haas, Emla Fitzsimons, and Heike Harmgart.** 2015. "The Impacts of Microfinance: Evidence from Joint-Liability Lending in Mongolia." *American Economic Journal: Applied Economics* 7 (1): 90–122.

- Augsburg, Britta, Ralph De Haas, Heike Harmgart, and Costas Meghir.** 2015. "The Impacts of Microcredit: Evidence from Bosnia and Herzegovina." *American Economic Journal: Applied Economics* 7 (1): 183–203.
- Banerjee, Abhijit, Emily Breza, Arun G. Chandrasekhar, and Benjamin Golub.** 2018. "When Less is More: Experimental Evidence on Information Delivery During India's Demonetization." 2018. NBER Working Paper 24679.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson.** 2013. "The Diffusion of Microfinance." *Science* 341 (6144).
- Banerjee, Abhijit, and Esther Duflo.** 2010. "Giving Credit Where It Is Due." *Journal of Economic Perspectives* 24 (3): 61–80.
- Banerjee, Abhijit, and Esther Duflo.** 2014. "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program." *Review of Economic Studies* 81 (2): 572–607.
- Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan.** 2015. "The Miracle of Microfinance? Evidence from a Randomized Evaluation." *American Economic Journal: Applied Economics* 7 (1): 22–53.
- Bass, Frank M.** 1969. "A New Product Growth for Model Consumer Durables." *Management Science* 15 (5): 215–27.
- Bertrand, Marianne, Antoinette Schoar, and David Thesmar.** 2007. "Banking Deregulation and Industry Structure: Evidence from the French Banking Reforms of 1985." *Journal of Finance* 62 (2): 597–628.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen.** 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4): 1347–93.
- Borusyak, Kirill, and Peter Hull.** 2022. "Non-Random Exposure to Exogenous Shocks: Theory and Applications." NBER Working Paper 27845.
- Breza, Emily, and Cynthia Kinnan.** 2021. "Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis." *Quarterly Journal of Economics* 136 (3): 1447–97.
- Brown, J. David, and John S. Earle.** 2017. "Finance and Growth at the Firm Level: Evidence from SBA Loans." *Journal of Finance* 72 (3): 1039–80.
- Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin.** 2021. "The Macroeconomics of Microfinance." *Review of Economic Studies* 88 (1): 126–61.
- Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel.** 2018. "Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets." *Quarterly Journal of Economics* 134 (2): 785–842.
- Cai, Jing, Alain de Janvry, and Elisabeth Sadoulet.** 2015. "Social Networks and the Decision to Insure." *American Economic Journal: Applied Economics* 7 (2): 81–108.
- Cai, Jing, and Adam Szeidl.** 2018. "Interfirm Relationships and Business Performance." *Quarterly Journal of Economics* 133 (3): 1229–82.
- Cai, Jing, and Adam Szeidl.** 2024. "Replication data for: Indirect Effects of Access to Finance." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E197302V1>.
- Calderon, Gabriela, Jesse M. Cunha, and Giacomo De Giorgi.** 2020. "Business Literacy and Development: Evidence from a Randomized Controlled Trial in Rural Mexico." *Economic Development and Cultural Change* 68 (2): 507–40.
- Chandrasekhar, Arun G., Benjamin Golub, and He Yang.** 2018. "Signaling, Shame, and Silence in Social Learning." NBER Working Paper 25169.
- Chodorow-Reich, Gabriel.** 2014. "The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008-9 Financial Crisis." *Quarterly Journal of Economics* 129 (1): 1–59.
- Crepon, Bruno, Florencia Devoto, Esther Duflo, and William Pariente.** 2015. "Estimating the Impact of Microcredit on Those who Take It Up: Evidence from a Randomized Experiment in Morocco." *American Economic Journal: Applied Economics* 7 (1): 123–50.
- Crepon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora.** 2013. "Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment." *Quarterly Journal of Economics* 128 (2): 531–80.
- De Mel, Suresh, David McKenzie, and Christopher Woodruff.** 2008. "Returns to Capital in Microenterprises: Evidence from a Field Experiment." *Quarterly Journal of Economics* 123 (4): 1329–72.
- Dolfen, Paul, Liran Einav, Peter J. Klenow, Benjamin Klopach, Jonathan D. Levin, Laurence Levin, and Wayne Best.** 2019. "Assessing the Gains from E-Commerce." NBER Working Paper 25610.
- Drexler, Alejandro, Greg Fischer, and Antoinette Schoar.** 2014. "Keeping It Simple: Financial Literacy and Rules of Thumb." *American Economic Journal: Applied Economics* 6 (2): 1–31.

- Duflo, Esther, and Emmanuel Saez.** 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 118 (3): 815–42.
- Einav, Liran, Peter J. Klenow, Jonathan D. Levin, and Raviv Murciano-Gorof.** 2021. "Customers and Retail Growth." Unpublished.
- Fafchamps, Marcel, David McKenzie, Simon Quinn, and Christopher Woodruff.** 2014. "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana." *Journal of Development Economics* 106: 211–26.
- Field, Erica, Rohini Pande, John Papp, and Natalia Rigol.** 2013. "Does the Classic Microfinance Model Discourage Entrepreneurship among the Poor? Experimental Evidence from India." *American Economic Review* 103 (6): 2196–2226.
- Fujita, Masahisa, and Jacques-François Thisse.** 1996. "Economics of Agglomeration." *Journal of the Japanese and International Economies* 10 (4): 339–78.
- Guiteras, Raymond, James Levinsohn, and Ahmed M. Mobarak.** 2019. "Demand Estimation with Strategic Complementarities: Sanitation in Bangladesh." Unpublished.
- Hardy, Morgan, and Jamie McCasland.** 2021. "It Takes Two: Experimental Evidence on the Determinants of Technology Diffusion." *Journal of Development Economics* 149: 102600.
- Hottman, Colin J., Stephen J. Redding, and David E. Weinstein.** 2016. "Quantifying the Sources of Firm Heterogeneity." *Quarterly Journal of Economics* 131 (3): 1291–1364.
- Huber, Kilian.** 2018. "Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties." *American Economic Review* 108 (3): 868–98.
- Jensen, Robert.** 2007. "The Digital Divide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector." *Quarterly Journal of Economics* 122 (3): 879–924.
- Leonardi, Marco, and Enrico Moretti.** 2022. "The Agglomeration of Urban Amenities: Evidence from Milan Restaurants." NBER Working Paper 29663.
- Marshall, Alfred.** 1920. *Principles of Economics*. 8th ed. London: Macmillan.
- McKenzie, David.** 2012. "Beyond Baseline and Follow-Up: The Case for More T in Experiments." *Journal of Development Economics* 99 (2): 210–21.
- McKenzie, David, and Susana Puerto.** 2021. "Growing Markets through Business Training for Female Entrepreneurs: A Market-Level Randomized Experiment in Kenya." *American Economic Journal: Applied Economics* 13 (2): 297–332.
- McKenzie, David, and Christopher Woodruff.** 2008. "Experimental Evidence on Returns to Capital and Access to Finance in Mexico." *World Bank Economic Review* 22 (3): 457–82.
- Ponticelli, Jacopo, and Leonardo S. Alencar.** 2016. "Court Enforcement, Bank Loans, and Firm Investment: Evidence from a Bankruptcy Reform in Brazil." *Quarterly Journal of Economics* 131 (3): 1365–1413.
- Romano, Joseph P., and Michael Wolf.** 2005. "Exact and Approximate Stepdown Methods for Multiple Hypothesis Testing." *Journal of the American Statistical Association* 100 (469): 94–108.
- Rotemberg, Martin.** 2019. "Equilibrium Effects of Firm Subsidies." *American Economic Review* 109 (10): 3475–3513.
- Sraer, David, and David Thesmar.** 2020. "How to Use Natural Experiments to Measure Misallocation." Unpublished.
- Tarozzi, Alessandro, Jaikishan Desai, and Kristin Johnson.** 2015. "The Impacts of Microcredit: Evidence from Ethiopia." *American Economic Journal: Applied Economics* 7 (1): 54–89.