

INTERFIRM RELATIONSHIPS AND BUSINESS PERFORMANCE*

JING CAI AND ADAM SZEIDL

We organized business associations for the owner-managers of young Chinese firms to study the effect of business networks on firm performance. We randomized 2,820 firms into small groups whose managers held monthly meetings for one year, and into a “no-meetings” control group. We find the following. (i) The meetings increased firm revenue by 8.1%, and also significantly increased profit, factors, inputs, the number of partners, borrowing, and a management score. (ii) These effects persisted one year after the conclusion of the meetings. (iii) Firms randomized to have better peers exhibited higher growth. We exploit additional interventions to document concrete channels. (iv) Managers shared exogenous business-relevant information, particularly when they were not competitors, showing that the meetings facilitated learning from peers. (v) Managers created more business partnerships in the regular than in other one-time meetings, showing that the meetings improved supplier-client matching. *JEL Codes:* D22, O12, O14, L14.

I. INTRODUCTION

Much research has focused on barriers to firm growth that act at the level of the individual firm, such as limits to borrowing or lack of managerial skills. But firms do not operate in a vacuum: business relationships, which provide information, training, referrals, intermediate inputs, and other services, are potentially central. Because of networking frictions such as lack of information or lack of trust, these relationships may not form efficiently,

*We thank Attila Gaspar, Huayu Xu, Hang Yu, and Zhengdong Zhang for excellent research assistance; Daron Acemoglu, Pol Antràs, David Atkin, Abhijit Banerjee, Andrew Bernard, Nick Bloom, Emily Breza, Arun Chandrasekhar, Esther Duflo, Ben Golub, Matt Jackson, Terence Johnson, Dean Karlan, Larry Katz, Sam Kortum, David Lam, Ben Olken, Rohini Pande, Mark Rosenweig, Antoinette Schoar, Matthew Shapiro, Duncan Thomas, Eric Verhoogen, Chris Woodruff, Dean Yang, and seminar participants for helpful comments. We thank the Innovations for Poverty Action’s SME Initiative, the Private Enterprise Development in Low-Income Countries (ERG 1893 and MRG 2355), the University of Michigan, the European Research Council under the European Union’s Seventh Framework Program (FP7/2007-2013) grant agreement number 283484, and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme grant agreement number 724501 for funding.

© The Author 2018. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

The Quarterly Journal of Economics (2018), 1229–1282. doi:10.1093/qje/qjx049.

Advance Access publication on December 20, 2017.

leading to a possible network-based growth barrier. Motivated by similar considerations, a small literature going back to [McMillan and Woodruff \(1999\)](#) has begun to explore the role of interfirm relationships for economic development.¹ But we still know little about the effect of an exogenous expansion of business networks on firm performance, the underlying mechanisms, and policies that can induce such a change.

We investigate these issues using a large-scale field experiment in China, in which we organized experimental business associations for the owner-managers of small and medium enterprises (SMEs). Building on existing approaches to induce variation in business connections—especially those by [Fafchamps and Quinn \(forthcoming\)](#) and [Bernard, Moxnes, and Saito \(forthcoming\)](#)—we created networks through regular meetings which had the explicit purpose of fostering business interactions. We also introduced additional interventions to learn about mechanisms. Our main findings are that business meetings substantially and persistently improved firm performance in many domains, and that learning and partnering were active mechanisms. These results suggest that differences in business networks may explain some of the large observed heterogeneity in firm performance ([Syverson 2011](#)). Because SMEs produce a large share of the output in developing countries, the results also suggest that organizing business associations can meaningfully contribute to private sector development.

In [Section II](#) we introduce our experimental design. In the summer of 2013 we invited micro and small and medium enterprises established in the preceding three years in Nanchang, China, to participate in business associations. From 2,820 firms that expressed interest, we randomly selected 1,500 and randomized their owner-managers into meetings groups with 10 managers each. We informed the remaining 1,320 firms—the control group—that there was no room for them in the meetings.

Managers in each meeting group were encouraged to hold monthly self-organized meetings. These meetings were intensive: in a typical meeting managers visited the firm of a group member, took a tour, and spent hours discussing business-relevant issues. The program lasted for one year. We surveyed the firms in summer 2013 before the intervention (baseline), in summer 2014 shortly after the end of the intervention (midline), and in summer

1. We review this literature in detail below.

2015 one year after the end of the intervention (endline). In the surveys we collected information on (i) Firm characteristics, including sales, employment, and other balance sheet variables; (ii) managerial characteristics, including—in the midline and endline surveys—management practices; and (iii) firm networks. As an incentive to participate in the intervention and the survey, we gave a certificate—providing access to certain government services—to complying treatment and control firms.

We introduced three additional interventions to document internal consistency and learn about mechanisms. First, to explore peer effects, we created variation in the composition of groups by sector and size. Second, to document learning, similarly to [Duflo and Saez \(2003\)](#) and [Cai, de Janvry, and Sadoulet \(2015\)](#), we provided randomly chosen managers with information about two financial products: a government grant and a private saving opportunity.² Third, to explore the role of meeting frequency, building on [Feigenberg, Field, and Pande \(2013\)](#) we organized one-time cross-group meetings for a subset of managers.

In [Section III](#) we present results on the effect of the meetings. We first explore the overall impact of the intervention. Our basic regression is a firm fixed effects specification that effectively compares the within-firm growth rate in the meetings groups to that in the control group. We estimate that by the midline survey the sales of treatment firms increased by a significant 7.8 log points more than that of control firms, corresponding to a treatment effect on sales of 8.1%. This effect persisted to the endline survey: the baseline-to-endline change in log sales was 9.8 points higher in treatment than in control firms ($p < .05$), corresponding to a long-term treatment effect on sales of 10.3%. We also find significant and persistent impacts for profits, production factors (employment and fixed assets), and inputs (materials and utility cost).

Turning to intermediate outcomes, we find that the meetings significantly and persistently increased the number of clients, the number of suppliers, and formal and informal borrowing. We also find that the meetings improved a management score—computed either from managers' or from workers' survey responses about business practices—by about 0.2 standard deviations ($p < .05$). A natural interpretation of this result, also supported by the fact that the management score predicts revenue conditional on

2. We also provided the information to random control firms to ensure that the same share of treatment and control firms were directly informed.

factors and inputs, is that the meetings increased firm productivity. We also find positive effects on innovation. Besides confirming the beneficial effects of the meetings, the results on intermediate outcomes suggest at least two possible underlying mechanisms: learning from peers, which may have improved management and innovation; and better firm-to-firm matching, which may have created new partnerships. But the results do not yet establish that these mechanisms were indeed active: it is possible that the meetings created growth through some other channel, which then led to an increase in intermediate outcomes.

We then explore the role of peer composition. We view this analysis as an internal consistency test that further supports our identification: plausible mechanisms operating through business networks all seem to imply that having better peers should improve performance. We proxy peers' quality with their size (employment) at baseline, and ask whether firms randomized into groups with larger peers grew faster. We find evidence for peer effects in several outcomes, including sales, profits, utility costs, the number of clients, and management practices. Overall these findings confirm, using a different source of variation, our basic result that business networks improve firm performance.

We next discuss some issues with identification and interpretation. One concern is that experimenter demand effects may drive the results. Contradicting this explanation, we find essentially no difference between the self-reported and the actual book value of sales. Demand effects are also unlikely to explain peer effects, which are identified using only firms in the treatment. Another concern is that the meetings may have had a side effect through improved access to government officials or the government grant opportunity about which we distributed information in an additional intervention. But access to government officials cannot easily explain peer effects and the gains in management and innovation. In addition, controlling for access to government funding does not change our results. Another side effect may be collusion: perhaps firms in the meetings coordinated price increases. However, standard models of collusion would predict a reduction in quantity, contradicting the positive effects on factors and inputs, and collusion cannot easily explain the gains in management.

In [Section IV](#) we document evidence for two mechanisms: learning and partnering. We begin with learning and show that the meetings diffused business-relevant information. We do this

using the additional intervention in which we provided information about two different financial products (independently) to randomly chosen managers. For both products, we find that uninformed managers in groups with informed peers were about 30 percentage points more likely to apply, providing direct evidence on learning as a mechanism. We also show that for the more rival product, a grant opportunity for the firm—which could help a competitor’s business—diffusion was weaker in groups in which firms on average had more competitors. In contrast, for the less rival product, a savings opportunity for the manager, diffusion was not weaker in groups with higher competition. These results suggest that the diffusion of rival information was limited by product market competition. In independent work, [Hardy and McCasland \(2016\)](#) show that the diffusion of a new weaving technique in Ghana was lower in treatments with higher experimentally induced competition. Taken together, their findings and ours highlight the potential relevance of an understudied friction in technology diffusion: the endogenous (dis)incentive to transmit information.

We document evidence on a second mechanism—improved access to partners—using the intervention of one-time cross-group meetings. We show that by midline firms received referrals from 2.2 more peers, and established direct partnerships—supplier, client, or joint venture—with 1.2 more peers in their regular group than in their cross-group ($p < .01$). These findings indicate that regular meetings reduced the cost of partnering. Differences in referral and partnership rates remained in the year after the conclusion of the meetings, showing that the intervention created persistent firm-to-firm connections. We also find that in hypothetical trust games managers exhibited significantly higher trust—at both midline and endline—towards their regular than their cross-group partners, suggesting a possible mechanism through which repeated interactions helped improve partnering.

In the concluding [Section V](#) we discuss several implications of the results. We begin with a cost-benefit calculation. A back-of-the-envelope estimate suggests that for the average firm the profit gains from the meetings were twice as high as the costs of organizing and attending. Thus the intervention appears to have been quite cost effective. A natural question is why managers did not organize meetings for themselves. There are several possible reasons. Search costs and trust barriers may be higher if managers have to organize the meetings themselves; there may be a

public good problem if these costs fall on a single organizer; and, paralleling the argument of [Bloom et al. \(2013\)](#), managers may have underestimated the gains from business associations.

We then compare our impacts to other interventions. Business training is often estimated to have modest and insignificant effects on firm performance ([McKenzie and Woodruff 2014](#)). For intensive and personalized management consulting [Bloom et al. \(2013\)](#) estimate a productivity gain of 17%. We find smaller effects—an 8% sales increase—but our intervention is cheaper and appears to be quite cost-effective. Their results and ours suggest that intensive interventions may have a higher chance of improving performance, perhaps through a “demonstration effect” of directly observing superior business practices. The fact that both their sample and our sample was selected suggests that interventions may have a larger effect when participants are interested in improving their business. We conclude that organizing regular business meetings for such firms can be an effective tool for private sector development.³

Our work builds on and contributes to three main literatures. Our research questions are most related to the work on firm-to-firm interactions. Theories in this area include [Acemoglu et al. \(2012\)](#), [Antràs and Chor \(2013\)](#), [Oberfield \(forthcoming\)](#), and [Eaton, Kortum, and Kramarz \(2015\)](#), who explore the aggregate and efficiency implications of supply chain networks. Evidence from observational data suggest that business networks can improve several firm outcomes, including access to credit ([McMillan and Woodruff 1999](#); [Khwaja, Mian, and Qamar 2011](#); [Haselmann, Schoenherr, and Vig forthcoming](#)), managerial compensation policy ([Shue 2013](#)), investment performance ([Hochberg, Ljungqvist, and Lu 2007](#)), and access to business partners ([Bernard, Moxnes, and Saito forthcoming](#); [Bernstein, Giroud, and Townsend 2016](#)).⁴ There is almost no experimental evidence on the impact of firm networks, except for the pioneering study by [Fafchamps and Quinn \(forthcoming\)](#), who document the diffusion of some management practices through connections created by joint committee membership. Our contribution to this literature is to

3. Sample selection and the demonstration effect may have also been important for the success of the management training trips organized under the Marshall Plan ([Giorcelli 2017](#)). Another example of a policy intervention broadly similar to ours but involving large firms and government agencies is the “Mesas ejecutivas” program in Peru ([Ministerio de la Produccion del Peru 2016](#)).

4. Also related is the work about agglomeration effects, reviewed in [Duranton and Puga \(2004\)](#) and [Rosenthal and Strange \(2004\)](#).

experimentally evaluate the impact of business networks on a broad range of firm outcomes and identify specific mechanisms.

Our methodology and policy results build on a literature that uses experiments to study private sector development. [De Mel, McKenzie, and Woodruff \(2008\)](#) measure the return to capital in microenterprises. Several papers reviewed in [McKenzie and Woodruff \(2014\)](#) study the effects of business training, while [Bloom et al. \(2013\)](#) and [Bruhn, Karlan, and Schoar \(forthcoming\)](#) measure the impact of management consulting. [McKenzie \(2017\)](#) evaluates a business plan competition and [Brooks, Donovan, and Johnson \(forthcoming\)](#) evaluate a business mentoring program. We contribute to this work with a large-scale experiment on the key but understudied segment of SMEs and by evaluating the new policy intervention of organizing business associations.

Our results on mechanisms relate to a literature on network effects in economics. This includes research on peer effects, information diffusion in networks, network-based referrals, and network-based trust.⁵ We contribute to this work by documenting peer effects, referrals, and the role of trust in the new domain of managerial networks, and especially to the work on information diffusion by documenting—together with [Hardy and McCasland \(2016\)](#)—the new mechanism that competition can limit the transmission of rival information.

II. CONTEXT, EXPERIMENTAL DESIGN, AND DATA

II.A. Context

Our experimental site was Nanchang, the capital city of Jiangxi Province, in southeastern China. In 2014 the city had a population of around 5 million people, and a GDP of \$58 billion, which ranked it as the 19th among the 32 capital cities in China. Nanchang was growing fast before the start of our study, with over 30,000 microenterprises and SMEs established during 2010–2013.

We conducted our intervention in collaboration with the Commission of Industry and Information Technology (CIIT) in Nanchang, one of the main government departments in charge of private sector development.

5. See, for example, [Sacerdote \(2001\)](#) on peer effects, [Banerjee et al. \(2013\)](#) on information diffusion, [Ioannides and Loury \(2004\)](#) on referrals, and [Karlan et al. \(2009\)](#) on trust. We review these literatures in more detail when we discuss the specific results below.




FIGURE I

Timeline and Interventions

II.B. Interventions

1. Basic Experiment. Figure I summarizes the timeline and interventions of our experiment. In summer 2013, through CIIT we invited all microenterprises and SMEs established in the preceding three years in Nanchang to participate in business associations. Around 5,400 firms expressed interest. We randomly selected 2,820 firms from this pool as our study sample. Almost all of these firms were owner managed, and from here on we refer to the CEO of the firm simply as the manager. Out of the study sample we randomly selected 1,500 firms—the treatment group—and randomized them into meetings groups with about 10 firms each.⁶ We informed the 1,320 control firms that there was no room for them in the meetings.

The managers in each meeting group were expected to meet once a month, every month, for one year. We organized the first meetings, in collaboration with CIIT, in August 2013. For this first meeting only, we offered the managers in each group print materials containing business-relevant information. We gave the same material to control firms as well. CIIT chose one of the managers in each meeting group to be the group leader. This person was responsible for planning and scheduling all subsequent monthly

6. To ensure that the managers of the firms in each meeting group were relatively close to each other, we divided the study area into 26 local subregions, and randomized firms into the treatment and control group, and treatment firms into meetings groups, at the subregion level.

meetings. Each meeting was attended by one of our surveyors, typically an undergraduate student at a local university, who took notes on the location, date, attendance, topics discussed, and the main takeaways and submitted the log to us.

According to the meeting logs, in most groups members took turns hosting the meetings. In a typical meeting, group members toured the firm of the host manager and then spent hours discussing business-relevant issues. Typical meetings lasted for about half a day. Common discussion topics included borrowing, management, suppliers and clients, hiring, recent government policies, and marketing. Average attendance in the meetings was 87%.

To provide incentives to participate, we offered a certificate from CIIT to managers in the control group who answered our surveys and managers in the treatment group who answered our surveys and attended at least 10 out of the 12 monthly meetings. The certificate stated that the firm was selected to be in the database of micro, small, and medium enterprises of Nanchang City.⁷ In China, “being selected into the database” can be a measure of excellence of individuals and organizations, such as experts, entrepreneurs, or companies. CIIT explained to managers in the invitation letter and at the baseline survey that being selected into the database allows for improved access to some of their services, including government funding and admission to local entrepreneur training programs. In addition, the certificate may also be viewed as a signal of firm quality. CIIT gave the certificate to firms in August 2014, after the conclusion of the one-year program and the midline survey. To get a direct measure of the certificate’s benefits, we asked all firms in the midline and endline surveys to report their subjective value—willingness to pay—for the certificate. As we discuss in more detail later, the average willingness to pay was not different between treatment and control firms and amounted to 0.7% of baseline profits or 0.04% of baseline sales.

2. Additional Interventions. To improve identification and explore mechanisms, we introduced three additional

7. The translation of the full text of the certificate is as follows. “You have been selected into the database of micro, small, and medium enterprises in Nanchang City, Jiangxi Province. Certificate issued by the Commission of Industry and Information Technology, August 2014.”

interventions. First, to help measure peer effects, we created variation in the composition of groups by size and sector. Almost all of our firms were from two sectors, manufacturing and services. In each subregion, we created two firm size categories, “small” and “large”, by the median employment of our sample of firms in that subregion. We then created four types of groups: (i) small firms in the same sector, (ii) large firms in the same sector, (iii) mixed-size firms in the same sector, and (iv) mixed size and mixed sector. We randomized treated firms into these groups in each subregion.

We implemented this randomization as follows. (1) In each of the 26 subregions we divided firms into four strata: (a) small service, (b) big service, (c) small manufacturing, and (d) big manufacturing. (2) In each strata of each subregion we randomly ranked firms. (3) In each subregion we created an assignment that mapped firms by their strata and rank into business groups of different types.⁸ (4) Using the random rankings, we implemented the assignments. The randomization ensured that conditional on the firm’s strata and subregion, the peers of the firm were random. We created the assignments ourselves, taking into account the number of firms of different types in a subregion and target values for the aggregate number of group types. Because CIIT staff expected that they would perform better, we targeted to have about 30% more mixed groups, and ended up with 30, 32, 40, and 43 type (i), (ii), (iii), and (iv) groups of firms.

In a second additional intervention, designed to measure information diffusion, we gave information about two relatively unknown financial products to randomly chosen managers. The first product was a funding opportunity for the firm: a government grant of up to RMB 200,000 (about US\$32,000 at that time) for which all firms in the region could apply. Because it could help the business of a competitor’s firm, managers may have viewed this

8. In each subregion the assignment was a collection of four vectors per strata. Each vector corresponded to a group type (e.g., small firms same sector) and the elements specified the number of firms to be assigned to each group of that type. The dimension of the vector measured the number of groups of that type. For example, in one subregion, strata (c) of 46 small manufacturing firms had the assignment vectors (11, 0), (0, 0), (5, 5, 5, 6), (4, 5, 5). Thus the first 11 and 0 firms from the strata were used for the the first and the second type (i) groups; no firms were used for type (ii) groups; the next 5, 5, 5, and 6 firms were used for the four type (iii) groups; and the next 4, 4, and 5 firms were used for the three type (iv) groups.

product to be rival and may have been unwilling to discuss it with competitors. The second product was a savings opportunity for the manager: a product offering an annual return of almost 7%, which was higher than the typical return of available high-yield saving products in the market (about 4–5%). Because it could not directly help a competitor's business, managers may have viewed this product to be less rival and may have been more willing to discuss it with competitors.⁹

We distributed information about each product in February 2014 via phone calls and text messages to 0%, 50%, or 80% of the managers in each meeting group. We randomly assigned about one third of the meeting groups to each of these three treatment intensities.¹⁰ We distributed the information to 40% of control firms to ensure that the same share of treatment and control firms have the information. We randomized and distributed the information independently for the two products.

The timing of the government grant was the following. The application period started in May 2014, and applications were due June 20, 2014, just before the midline survey, which took place in July and August. Decisions were made in December 2014, and the grants were paid out in February 2015. Thus the grant itself could not have directly affected firm outcomes at midline, although anticipation effects may have played a role. We discuss these issues in more detail later.

As a final intervention, to learn about the role of meeting frequency, we organized one-time cross-group meetings. We randomized 466 managers in the meetings treatment into 43 “cross-groups” of about 10 managers each, such that no two managers from the same meetings group were in the same cross-group.¹¹ Each cross-group met once, in February 2014.

II.C. Surveys

We conducted a baseline survey before the intervention in summer 2013, a midline survey after the intervention in summer

9. Both products were in limited supply.

10. We stratified this randomization by group type.

11. The basic logic of the randomization was the following. We randomly selected 80% of the regular groups in each of the four group types. We randomly picked four firms in each selected group. Then at the subregion level we sequentially randomly assigned these firms into cross-groups, ensuring that to any given cross-group at most one firm is assigned from each regular group.

2014, and an endline survey in summer 2015. Because the fiscal year in China ends in June, data in the baseline survey refer to the fiscal year before the intervention; data in the midline survey refer to the fiscal year that almost fully overlaps with the meetings; and data in the endline survey refer to the fiscal year after the conclusion of the meetings.

The surveys were conducted in person with the manager, by locally hired enumerators, in collaboration with CIIT. CIIT officials phoned the firms in advance to arrange the interview, and if the manager was not available at the scheduled time, a CIIT official or our enumerator phoned again to arrange a second meeting. In addition, a CIIT official was present at each interview to help build trust between the manager and our enumerator.

In the surveys we collected information from both treatment and control firms about the following groups of variables.

(i) Firm characteristics: profits, sales, costs, utility expenses, spending on intermediate inputs, and other balance-sheet variables. For sales we have two measures: besides the self-reported value in the survey, we have the actual book value. To obtain it, at the conclusion of the survey our enumerators asked the accountant of the firm to physically show the value in the firm's book.¹² (ii) Managerial characteristics: demographics, measures of well-being, and—in the midline and endline survey—questions on management. (iii) Firm networks: the number and type of business connections (supplier, buyer, joint venture) within and outside the group and information on the nature of any relationship with group members (competitor or some type of partner).¹³ (iv) Whether managers applied for the funding opportunities about which we had distributed information. (v) Other outcomes: these included product innovation, and also, for a random subset of 750 firms, a survey of one randomly picked worker per firm on working conditions. We only included these other outcomes in the endline survey. The English version of our survey questions is available in Section O3 of the [Online Appendix](#).

Many of the areas of firm behavior above are commonly surveyed, and accordingly we mostly relied on standard

12. This procedure worked for most firms. When the firm did not have an accountant or the accountant was not present, we asked the manager to show us the book.

13. Because the firms in each group came from a large pool, there were essentially no preexisting in-group partnerships at baseline.

questions. The key novelty was in the area of management practices, where—building on the pioneering work of Bloom and Van Reenen (2007)—we developed a questionnaire suited for our sample of SMEs in China. Our starting point was the 2010 Manufacturing Survey Instrument of the World Management Survey (WMS). The WMS is administered using open-ended questions by specifically trained surveyors, a technique we were unable to implement given subjects' time constraints and our resource constraints. Similarly to Bruhn, Karlan, and Schoar (forthcoming), we thus opted for asking the managers directly about concrete management practices. Because our sample consisted of smaller and less developed firms than those commonly included in the WMS, we modified their survey by omitting some questions, simplifying others, and adding more basic questions. For example, some questions in the WMS ask about lean (modern) management techniques. As managers in our sample were unlikely to be sufficiently familiar with the notion of lean management techniques, we omitted those questions. Other questions in the WMS survey ask about performance tracking and key performance indicators (KPIs), starting with "What kind of KPIs would you use for performance tracking?" We simplified these questions by focusing more narrowly on employee performance and asking questions such as "On average, how often do you evaluate the performance of your employees? (months)" and "Do you track employee performance using numerical performance indicators (e.g., number of items sold)? (1 = yes, 0 = no)" We piloted our management questions with a sample of about 100 firms and made adjustments to ensure that managers found them clear and relevant.

Our final management survey consisted of 19 questions and covered five areas of management: evaluation and communication of employee performance, targets and responsibilities, attracting and incentivizing talent, process documentation and development, and delegation. Below we show evidence that these data contain information both about firm performance and about employees' perceptions of management practices.

II.D. Summary Statistics and Randomization Checks

Table I shows basic summary statistics from the baseline survey. The first three columns report the means for all firms, treatment firms, and control firms; the final column reports the difference between treatment and control firms. Panel A on firm

TABLE I
SUMMARY STATISTICS: FIRM AND MANAGER CHARACTERISTICS

	All sample	Treatment	Control	Difference
Number of Observations	2,820	1,500	1,320	
Panel A: Firm characteristics (2013 baseline)				
Firm age	2.34 (1.75)	2.39 (1.72)	2.29 (1.77)	0.10 (0.07)
Ownership: domestic private firms	0.98 (0.15)	0.98 (0.15)	0.98 (0.15)	0.00 (0.01)
Sector: manufacturing	0.50 (0.50)	0.51 (0.50)	0.48 (0.50)	0.03 (0.02)
Sector: service	0.48 (0.50)	0.47 (0.50)	0.49 (0.50)	-0.02 (0.02)
Number of employees	36.19 (86.49)	36.33 (90.63)	36.01 (81.55)	0.32 (3.37)
Panel B: Managerial characteristics (2013 baseline)				
Gender (1=male, 0=female)	0.84 (0.37)	0.85 (0.36)	0.84 (0.37)	0.01 (0.01)
Age	40.84 (8.85)	41.05 (8.46)	40.59 (9.27)	0.46 (0.34)
Education: college	0.29 (0.45)	0.29 (0.45)	0.30 (0.46)	-0.01 (0.02)
Government working experience	0.23 (0.42)	0.24 (0.42)	0.22 (0.41)	0.02 (0.02)
Communist Party member (1=yes, 0=no)	0.21 (0.4)	0.21 (0.4)	0.20 (0.4)	0.01 (0.02)

Notes. Standard deviations in parentheses for the first three columns. The fourth column reports the difference in characteristics between the treatment and control groups, standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

characteristics shows that in 2013 average firm age was about 2.3 years and that 98% of firms were domestic private enterprises.¹⁴ About half of the firms were in manufacturing and 48% in services.¹⁵ Consistent with self-selection of better firms into our sample, in spite of their young age these firms employed on average 36 workers. But the large standard deviation of employment (86) shows that there was much cross-firm heterogeneity.

Panel B presents managerial characteristics. The vast majority of managers were men, and in 2013 they were on average 41 years old. Almost a third of them had a college degree. Many

14. The remaining 2% were either privatized formerly state-owned firms, whose CEOs were appointed by the government, or foreign-owned firms. In both cases the local CEO was responsible for essentially all business-relevant decisions and is the person we label the manager.

15. Among others, firms in the manufacturing sector included textile, automobile, and furniture companies; and firms in the service sector included restaurants, wholesalers, and transportation companies.

TABLE II
SUMMARY STATISTICS: BUSINESS ACTIVITIES

	All sample	Treatment	Control	Difference
Number of Observations	2,820	1,500	1,320	
Panel A: Partnership (2013 baseline)				
Number of clients	45.89 (57.37)	45.58 (56.16)	46.23 (58.74)	-0.65 (2.24)
Number of suppliers	16.38 (19.23)	16.70 (20.30)	16.02 (17.94)	0.68 (0.75)
Panel B: Borrowing (2013 baseline)				
Bank loan (1=yes, 0=no)	0.25 (0.43)	0.25 (0.44)	0.25 (0.43)	0.00 (0.02)
Informal loan (1=yes, 0=no)	0.12 (0.33)	0.11 (0.32)	0.13 (0.34)	-0.02 (0.013)
Panel C: Accounting (2013 baseline)				
Sales (10,000 RMB)	1,592.70 (6,475.18)	1,510.62 (5,291.86)	1,686.19 (7,603.11)	-175.57 (252.32)
Log sales	5.59 (2.01)	5.61 (1.99)	5.58 (2.02)	0.03 (0.08)
Net profit (10,000 RMB)	79.23 (205.35)	77.26 (199.92)	81.52 (211.55)	-4.25 (8.09)
Panel D: Attrition and shutdown (relative to baseline sample)				
Attrition (2014 midline, %)	6.21 (24.13)	6.33 (24.36)	6.06 (23.87)	0.27 (0.91)
Attrition (2015 endline, %)	9.08 (28.73)	9.27 (29.01)	8.86 (28.43)	0.41 (1.08)
Shutdown (2015 endline, %)	10.25 (30.33)	10.20 (30.27)	10.30 (30.41)	-0.10 (1.14)
Panel E: Valuation of the CIIT certificate				
2014 midline (10,000 RMB)	0.56 (0.25)	0.56 (0.25)	0.56 (0.26)	-0.00 (0.01)
2015 endline (10,000 RMB)	0.56 (0.26)	0.56 (0.26)	0.56 (0.26)	-0.00 (0.01)

Notes. Standard deviations in parentheses for the first three columns. The fourth column reports the difference in characteristics between the treatment and control groups, standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

managers had government connections: 23% had worked either in government or in state-owned firms, and 21% of them were members of the Communist Party of China. There are no significant differences between the treatment and control firms in any of the variables in the table, confirming that our randomization is valid.

Table II shows summary statistics on firms' business activities. Panels A and B present data on business connections with suppliers, clients, and lenders. The average firm seems to have had a substantial customer and supplier base, with 46 clients and 16 suppliers. About 25% of firms borrowed from formal banks and

12% borrowed from friends and relatives in the previous year. The relatively large share of informal borrowing suggests frictions in getting formal loans, perhaps because they often require collateral or government guarantors.

Panel C reports data on accounting measures of firm performance. The average net profit was RMB 792,300 (about \$130,000), but this masks a lot of heterogeneity as indicated by the large standard deviation. A unitless measure of heterogeneity is the coefficient of variation (standard deviation divided by the mean), which for log sales is 0.36, higher than but roughly comparable to the corresponding value of 0.26 in the [Banerjee and Duflo \(2014\)](#) administrative data on mid-sized Indian firms. Consistent with the randomization, there are no significant differences between treatment and control firms in any of these variables.

Panel D reports measures of attrition and shutdown. Attrition is defined as one in a survey wave if we do not have information about the firm in that wave. Attrition can be the result of the firm choosing not to respond, moving away, or shutting down. We made a considerable effort to keep attrition low. With the help of CIIT we were able to track most mover firms; CIIT phoned managers in advance to arrange the survey, and when the manager was unavailable at the arranged time, we attempted to arrange a second meeting. The table shows that the attrition rates—relative to the baseline sample—at midline (about 6.21%) and at endline (about 9.08%) were not significantly different between treatment and control firms. In [Appendix Table A.1](#) we show that the baseline characteristics of attriting firms were also not significantly different between treatment and control firms. These facts indicate that selective attrition is unlikely to bias our results.

Panel D also reports the share of firms which we classify—based on the survey or direct information from CIIT—to have shut down by the endline survey. These firms constitute neither a subset nor a superset of the set of attriting firms. For some attriting firms we do not have information on the termination of operations, and these we do not classify as shutting down. Conversely, some firms that shut down still reported data for the part of the year during which they were active, and these we do not classify as attriting. The shutdown rate was about 10.25% for the full sample, and was slightly but insignificantly lower for treatment firms.

Panel E shows how much subjects valued the CIIT certificate, used as the incentive to participate. The average willingness to pay for it was about RMB 5,600 at both midline and endline (we did

not include this question at baseline), not significantly different between treatment and control firms. Thus differential access to government services via the use of the certificate is unlikely to have been an active force in our setting. The subjective value of the certificate amounted to 0.7% of baseline average profits or 0.04% of baseline average sales, suggesting that it was viewed to be valuable, but not so valuable that it would interfere with firm operations in a substantive way.

Because our sample consists of firms that responded to the invitation to participate in business associations, it is not representative. To get a sense of selection, we conducted a short survey of 124 randomly chosen nonresponding firms from the pool we originally contacted—microenterprises and SMEs in Nanchang created during a three-year window before summer 2013. [Appendix Table A.2](#) shows the results. As expected, nonresponding firms were smaller: on average they had half as many employees and a third as high revenues and profits as firms in our sample. They were also somewhat less likely to be run by a manager who is male or a member of the Communist Party. Due to this self-selection, our treatment effect estimates apply not to the representative firm, but to firms interested in participating in business associations. Importantly, because the treatment was introduced after the self-selection stage, our effects are identified for this sample. We think that the initial self-selection is a strength of our design, because it allows us to focus on a key segment of firms: those interested in improving themselves. These firms are relevant for economics because they are more likely to become successful and relevant for policy because they are the ones who respond to a policy intervention.

III. BUSINESS MEETINGS AND FIRM PERFORMANCE

In this section we show that the meetings improved firm performance in many domains, and also that firms randomized into groups with better peers grew faster. In the next section we study mechanisms.

III.A. *Effect of Meetings*

1. *Graphical Evidence.* We begin the analysis with graphical evidence that highlights some key patterns in the data. [Figure II](#) plots the kernel density of log sales for the treatment and the control group at baseline and at endline. Given that the surveys




FIGURE II
Kernel Density of Log Sales

were conducted at fiscal year end, the baseline data refer to the 12-month period before the start, and the endline data refer to the 12-month period after the end of the one-year meetings intervention. The left panel shows that—consistent with the randomization—before the intervention the distribution of log sales was similar in the treatment and control groups. The right panel shows that one year after the intervention the distribution of log sales for treatment firms was—slightly but visibly—to the right of that for control firms. The shift is present for a large part of the domain, showing that the meetings treatment increased sales for a substantial range of firm sizes. Although the shift seems visually small, this is mainly because the large heterogeneity of log sales leads to a wide range on the horizontal axis in the figure.

To quantify the shift and explore other outcomes, we turn to regressions.

2. *Empirical Strategy.* Our main empirical specification is

$$\begin{aligned}
 y_{it} = & \text{const} + \beta_1 \cdot \text{Midline}_{it} + \beta_2 \cdot \text{Endline}_{it} \\
 & + \beta_3 \cdot \text{Meetings}_{it} \times \text{Midline}_{it} + \beta_4 \cdot \text{Meetings}_{it} \times \text{Endline}_{it} \\
 (1) \quad & + \text{Firm f.e.} + \varepsilon_{it}.
 \end{aligned}$$

Here i indexes firms, t indexes years, and y_{it} is an outcome variable such as log sales. Meetings_{it} is an indicator for the treatment, which is time-invariant and equals 1 if the firm is invited to the meetings. Midline_{it} is an indicator for the midline survey wave, and Endline_{it} is an indicator for the endline survey wave. The firm fixed effects take out time-invariant heterogeneity, including

whether the firm is in the meetings treatment or in the control group. This specification is analogous to the one used by [De Mel, McKenzie, and Woodruff \(2008\)](#).

Our coefficients of interest are β_3 and β_4 , which measure—given the fixed effects specification—the differential change over time in the outcome variable in the treatment group relative to that in the control group. Intuitively, β_3 is the treatment-induced additional growth in y between baseline and midline; β_4 is the treatment-induced additional growth in y between baseline and endline. These coefficients can be compared to β_1 and β_2 , which measure the growth in y for the firms in the control group. The key identification assumption is that firms in the treatment group did not have systematically different trajectories from those in the control group for reasons other than the meetings treatment itself. Because the treatment is randomized, any potential omitted variable would have to be a side effect of the treatment itself, such as better access to government officials. We discuss possible omitted variables as we present the results and in [Section III.C](#). Because the treatment can induce correlated errors within a group, for inference we cluster standard errors at the level of the meeting group for treatment firms and at the level of the firm for control firms.

Our main sample includes all firms in all survey waves in which they responded, for a total of 7,857 observations. Due to attrition over time, this sample is an unbalanced panel, but as discussed in [Section II.D](#) attrition rates and attriting firms were not significantly different between treatment and control firms.¹⁶ To control for potential outliers, in specifications in which it is not binary, standardized, or bounded between 0 and 1, we winsorize the dependent variable at 1% in both tails of the distribution.¹⁷

3. Results. [Table III](#) presents results for a range of firm performance measures. Start with column (1) where the outcome is log sales, and consider first the effect at midline, that is, the fiscal year in which the meetings took place. While log sales in the control group increased, from baseline, by an insignificant 0.004,

16. Because not all firms responded to all survey questions, there are small reductions in sample size for some outcomes; response rates were not significantly different between treatment and control for any of them.

17. We show in [Table O1](#) of the [Online Appendix](#) that nonwinsorized specifications yield similar results.

TABLE III
EFFECT OF MEETINGS ON FIRM PERFORMANCE

Dependent var.:	log Sales (1)	Profit (10,000 RMB) (2)	log Number of employees (3)	log Total assets (4)	log Material cost (5)	log Utility cost (6)	log Productivity (7)
Midline	0.004 (0.019)	11.886** (5.402)	0.018 (0.017)	0.013 (0.017)	0.0003 (0.023)	-0.022 (0.021)	-0.010 (0.010)
Endline	0.013 (0.029)	12.213 (8.278)	0.029 (0.024)	0.019 (0.031)	0.023 (0.029)	0.024 (0.027)	0.007 (0.016)
Meetings*midline	0.078** (0.036)	25.746** (12.587)	0.052** (0.026)	0.061** (0.031)	0.055 (0.041)	0.099*** (0.036)	0.037** (0.017)
Meetings*endline	0.098** (0.049)	32.596* (18.525)	0.077* (0.044)	0.104** (0.047)	0.091* (0.054)	0.116** (0.046)	0.025 (0.025)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7,857	7,664	7,857	7,857	7,857	7,676	7,857
Mean dep. var. for control firms	5.587	104.259	2.706	3.959	4.882	1.831	1.590

Notes. Standard errors clustered at the meeting group level for treated firms and at the firm level for control firms. *** $p < .01$, ** $p < .05$, * $p < 0.1$.

log sales in the meetings treatment increased by an additional significant 0.078, corresponding to a treatment effect on sales of 8.1%. This effect persisted in the fiscal year after the meetings program ended: the coefficient of the interaction between *Meetings* and *Endline* shows that sales growth between baseline and endline was 9.8 log points higher for treated than for control firms, corresponding to a 10.3% treatment effect on sales. Similarly, column (2) shows that average profits increased by a significant RMB 257,500 (about \$36,000) more in the treatment group than in the control group by midline, and the difference persisted by endline. These results show large impacts for two key business relevant outcomes.

The remaining columns look at various components of the production process. Columns (3) and (4) show evidence on factors. We estimate significant and persistent treatment effects on both employment and fixed assets, ranging from 5 to 11 log points. Columns (5) and (6) focus on intermediate inputs. The treatment effect on materials is an insignificant but positive 5.5 log points by midline, which increases further to a significant 9.1 log points by endline ($p < .1$). The treatment effect on the utility cost is positive and highly significant throughout. Finally, column (7) shows the impact on total factor productivity, which we inferred using coefficients from estimating a revenue production function in the control group.¹⁸ The effect is only significant at midline. We do not read much into this result, because it is imprecise and subject to the identification problems associated with estimating production functions using revenue data (De Loecker 2011). To avoid those problems, below we focus on management, which we interpret as a component of productivity that we can measure more directly.¹⁹ Overall we conclude that Table III shows large and persistent benefits from the meetings.

Table IV explores intermediate outcomes that may have contributed to firm growth, as well as some alternative explanations. Columns (1) and (2) show highly significant and persistent treatment effects on the number of clients and suppliers, ranging

18. We inferred the coefficients from control firms to avoid the treatment confounding our production function estimate. The alternative approach of regressing log sales on the treatment as in equation (1) while controlling for factors and inputs yields almost identical estimates.

19. Also note that a 3 log point productivity gain could generate the observed growth in sales and factors under a demand elasticity of 3, which is well within the ballpark of standard estimates (Hsieh and Klenow 2009).

TABLE IV
INTERMEDIATE OUTCOMES AND ALTERNATIVE EXPLANATIONS

Dependent var.:	log Number of clients (1)	log Number of suppliers (2)	Bank loan (3)	Innovation (4)	log Reported - log book sales (5)	Tax/sales (6)
Midline	0.015 (0.020)	0.027 (0.021)	-0.040*** (0.011)		-0.001 (0.007)	0.001 (0.001)
Endline	0.044 (0.029)	0.049* (0.029)	0.008 (0.014)		-0.007 (0.006)	0.0017 (0.0012)
Meetings*midline	0.090*** (0.030)	0.085*** (0.031)	0.091*** (0.016)		-0.001 (0.011)	0.001 (0.001)
Meetings*endline	0.118** (0.046)	0.090** (0.041)	0.079*** (0.019)	0.082*** (0.028)	-0.002 (0.009)	-0.002 (0.002)
Firm fixed effects	Yes	Yes	Yes	No	Yes	Yes
Firm demographics	No	No	No	Yes	No	No
Observations	7,841	7,826	7,857	2,646	7,796	7,849
Mean dep. var. for control firms	3.211	2.13	0.239	0.123	0.028	0.024

Notes: Standard errors clustered at the meeting group level for treated firms and at the firm level for control firms. Firm demographics are indicators for firm size (above median employment in subregion at baseline), sector, subregion, and their interactions. *** $p < .01$, ** $p < .05$, * $p < .1$.

between 8 and 12 log points. Column (3) shows that firms in the meetings treatment were significantly more likely to take out loans following the intervention. For simplicity we group formal and informal loans into one indicator, but separately estimating treatment effects shows significant gains for both of them. These results can be interpreted in two ways. One possibility is that the meetings helped firms connect with more business partners and raise more capital, which contributed to firm growth. An alternative is that the meetings generated growth through other mechanisms, which translated into higher demand for business partners and capital. In [Section IV](#) we show direct evidence that improved partnering was one benefit of the meetings.

Column (4) shows the treatment effect on innovation, defined as an indicator for whether the firm introduced new products or services in that fiscal year. Because we asked about innovation only in the endline survey, we estimate a regression without firm fixed effects:

$$(2) \quad y_i = \text{const} + \beta_4 \cdot \text{Meetings}_{it} \times \text{Endline}_{it} + \text{Firm controls} + \varepsilon_i.$$

Because this regression only uses data from the endline survey, replacing the interaction with the uninteracted *Meetings* variable would yield the same coefficient β_4 . We report it as the coefficient of an interaction only to maintain the consistency of the table. Instead of firm fixed effects we control for a set of firm demographics: indicators for the firm's subregion, size category (above or below the median employment in our sample in the subregion), sector (manufacturing or services), and all their interactions. These are our standard set of firm controls used in several specifications in the article. Because the treatment is randomized, even in the absence of firm fixed effects β_4 continues to be identified: it reflects the difference in the level (not the growth rate) of innovation between the treatment and the control group. Because the purpose of innovation is to increase output given inputs, the significant positive estimate of 8.2 percentage points may represent future productivity gains due to the meetings.

Columns (5) and (6) focus on particular alternative explanations. Column (5) reports the treatment effect on the difference between the log of self-reported sales and the log of the book value of sales (which our enumerators took directly from the firm's book). There is no treatment effect on this difference, suggesting that experimenter demand effects are unlikely to drive the main results.

TABLE V
EFFECT OF MEETINGS ON FIRM MANAGEMENT

Dependent var.:	Management score (standardized)					
	Overall (1)	Evaluation (2)	Target (3)	Incentive (4)	Operation (5)	Delegation (6)
Meetings*midline	0.211*** (0.051)	0.094** (0.046)	0.034 (0.043)	0.237*** (0.047)	0.159*** (0.05)	0.071* (0.041)
Meetings*endline	0.215*** (0.048)	0.096** (0.045)	0.021 (0.046)	0.223*** (0.047)	0.179*** (0.044)	0.070 (0.043)
Observations	5,211	5,211	5,211	5,211	5,211	5,211
Mid/endline*firm demographics	Yes	Yes	Yes	Yes	Yes	Yes

Notes. Standard errors are clustered at the meeting group level for treated firms and at the firm level for control firms. Column (1) reports the impact of the treatment on the overall management z -score. Columns (2)–(6) report the impact on five components of management: evaluation and communication of employee performance; targets and responsibilities; attracting and incentivizing talent; process documentation and development; and delegation. *** $p < .01$, ** $p < .05$, * $p < .1$.

Column (6) shows that the tax-to-sales ratio of both treatment and control firms was essentially unchanged after the intervention. Thus improvement in tax avoidance is unlikely to have been the channel of the treatment effect.

4. *Management.* We turn to the effect of the treatment on management practices. Following Bloom and Van Reenen (2007) we aggregate the responses to management questions into a single index by standardizing, averaging, and standardizing them again. Because only the follow-up surveys contain data on management, we estimate an analogous specification to the one we used for innovation, which does not include firm fixed effects (but is still causally identified):

$$(3) \quad y_i = \text{const} + \beta_2 \cdot \text{Endline}_{it} + \beta_3 \cdot \text{Meetings}_{it} \times \text{Midline}_{it} + \beta_4 \cdot \text{Meetings}_{it} \times \text{Endline}_{it} + \text{Firm controls} + \varepsilon_i.$$

Table V reports the results. In column (1), we estimate treatment effects of 0.21 at both midline and endline ($p < .01$), measured in units of the standard deviation of the overall management score. In columns (2)–(6) we look at the treatment effect on different areas of management. We find that the intervention improved four of the five areas of management we surveyed, the exception being transparency of targets and responsibilities to

employees. Overall, we conclude that the meetings treatment had a large and persistent positive effect on management practices.

Given the argument in Bloom and Van Reenen (2007) and Bloom et al. (2013) that management is a component of total factor productivity, a natural interpretation of the results is productivity gains from the meetings. To strengthen this interpretation we present additional evidence that exploits a more direct measure of management practices and directly links our management score to firm performance. This evidence also addresses the concern that our management score—based on survey questions rather than the in-depth interviews of Bloom and Van Reenen (2007)—may just reflect improved use of business language, or the realization of the importance of management, but not the implementation of improved practices.

Our additional evidence exploits data on management that comes from a different source: the worker survey we conducted at the endline, with one worker each in a random subset of 739 sample firms (433 treatment, 306 control).²⁰ Because this survey asks the workers who are affected, it provides a more direct measure of actual business practices. Specifically, workers were surveyed about human resources (HR) practices in the following domains: evaluation and communication of employee performance, incentivizing talent, and delegation. Multiple questions closely corresponded to questions in our main management survey.²¹ Using the responses from the worker survey we constructed a standardized score that measures HR management (including delegation) from the perspective of workers.

The first column of Table VI shows the treatment effect of the meetings on this HR management score. Because the treatment is randomized, this regression is identified, and shows that the meetings improved HR practices as perceived by workers by 0.21 of a standard deviation ($p < .05$). In column (2) we regress the HR management score reported by workers on the HR component of our main management score variable, constructed by averaging and standardizing the manager's responses in the three HR-related domains. The significant coefficient of 0.13 shows that the two

20. We approached 750 firms for the worker survey, 11 of them did not provide answers (7 treatment, 4 control).

21. For example, in our main management survey we asked the manager "After the evaluation do you tell employees how they performed? (1 = Yes; 0 = No)", while in the worker survey we asked the worker whether "The company communicates with employees on how they performed after each evaluation. (1 = Strongly disagree; 2 = Disagree; 3 = Neutral; 4 = Agree; 5 = Strongly agree)."

TABLE VI
HR PRACTICES, MANAGEMENT SCORE, AND PRODUCTIVITY

Variables	HR management score reported by worker (standardized)			log Sales
	(1)	(2)	(3)	(4)
Meetings	0.208** (0.082)			
Management score (HR areas, standardized)		0.129*** (0.0378)	0.127*** (0.039)	
Management score (All areas, standardized)				0.020* (0.012)
log Total assets			0.039 (0.022)	0.063*** (0.019)
log Number of employees			0.012 (0.047)	0.093*** (0.016)
log Material cost			0.009 (0.025)	0.641*** (0.027)
Firm fixed effects	No	No	No	Yes
Firm demographics	Yes	Yes	Yes	No
Observations	739	739	725	5,147

Notes. Standard errors are clustered at the meeting group level for treated firms and at the firm level for control firms. Columns (1)–(3) use endline data for the subsample of firms for which the worker survey was conducted. Column (4) uses data in the midline and endline surveys. Firm demographics are firm size category, sector, subregion, and their interactions. *** $p < .01$, ** $p < .05$, * $p < .1$.

different measures of HR management share a common component. In column (3) we enrich this specification by including production factors and inputs; the unchanging estimate shows that the association is not driven by firm size. Finally, in column (4) we estimate a revenue production function using the full sample in the midline and endline survey waves, with firm fixed effects, in which we also include our main management score as an explanatory variable. Its loading of 0.02 ($p < .1$), shows that even controlling for factors, inputs, and firm-specific characteristics, variation in the management score is associated with variation in revenue.

In summary, [Table VI](#) provides more direct evidence that the meetings improved business practices, validates our management score as a measure of these practices, and shows that this score contains relevant information about firm performance. These results support our interpretation that the meetings improved firm productivity.

5. *Heterogeneous Effects.* As a final exercise in our analysis of the main effects, in Section O1 of the [Online Appendix](#) we report heterogeneous effect estimates of the meetings. The main result is that larger firms benefited more. We do not find heterogeneity along other firm and managerial characteristics. We also report treatment effects estimated separately for the four group types. We find generally positive effects—though less significant because of reduced power—in three group types: large same-sector firms, mixed-size same-sector firms, and mixed-size mixed-sector firms. The exception is the group type with small same-sector firms. These patterns are consistent with the finding on heterogeneity by firm size. The size heterogeneity result may also explain why the shutdown rate was not significantly different between treatment and control firms: treatment effects were smaller for precisely those firms—in the left tail of the size distribution—among which exit is more likely.

In summary, the results in [Tables III–VI](#) show that the meetings treatment substantially improved firm performance on several margins. The results on innovation and especially on management suggest genuine productivity gains. The effects on intermediate outcomes, taken together, suggest at least two mechanisms at work: learning from peers, which may have improved management, and improved partnering, which may have increased the number of suppliers and clients.

III.B. Group Composition and Peer Effects

We turn to estimate peer effects: whether being grouped with better peers at baseline improves a firm's performance. We view this analysis as an internal consistency test, because any mechanism we can think of that represents genuine network-based gains also predicts that the quality of peers should matter. Motivated by models such as [Melitz \(2003\)](#) in which productivity determines firm size, in our basic specification we measure peer quality with peer size (employment) at baseline. Using only the sample of firms in the meetings groups, our starting point is the following specification:

$$y_{it} = \text{const} + \delta_1 \cdot \text{Post}_{it} + \delta_2 \cdot \text{Post}_{it} \times \log \text{Peer size}_{it} + \text{Controls} \\ (4) \quad + \text{Firm f.e.} + \varepsilon_{it}.$$

Here $\log \text{Peer size}_{it}$ is the average of log employment of the other firms in the meeting group of firm i at baseline, that is, in the year before the intervention. Post_{it} is an indicator for both the midline and the endline survey waves: to increase power, we do not separate out peer effects by wave. The controls include the interaction of Post_{it} with our standard set of firm demographics, and are described in detail below.

The coefficient of interest in this regression is δ_2 , and we expect it to be positive because having better peers should improve performance. Importantly, $\delta_2 > 0$ should only be interpreted as evidence for composition effects, but not evidence that peer size creates improved performance. Indeed, peer size is likely to be correlated with several peer characteristics such as managerial skills, supply chains, and others, each of which may directly contribute to peer effects.

The main issue with consistently estimating δ_2 in our setting is that, as described in Section II.B, the group assignment was randomized only conditional on the subregion and strata of the firm. For example, in subregions in which average firm size was larger, firms mechanically tended to have larger peers. Because this variation in peer size is not random, it should not be used to identify δ_2 . We address this problem by including in the controls the interaction of Post_{it} with all the variables on which the random assignment of firms into groups was conditioned: indicators for subregion, sector categories at baseline (manufacturing or services), size categories at baseline (above or below median employment in the subregion), and all their interactions. With these controls, δ_2 is identified only from the random component of group assignment, and we report specification checks that make this transparent.

Table VII shows the results from estimating equation (4). Panel A shows the peer effect coefficients for our main firm performance measures. Column (1) implies that firms randomized into groups with 10 log points larger peers experienced an additional (significant) sales increase of 1.05 log points as a result of the treatment. That is, roughly, having 10% larger peers increased firm sales by 1%. For comparison, (log) peer size has an unconditional standard deviation of 0.9 and a conditional standard deviation—after controlling for the firm demographics on which the randomization was conditioned—of 0.49. Column (2) shows that 10 log points larger peers also increased annual firm profits by a significant RMB 27,825 (about \$4,500 at that time).

TABLE VII
EFFECT OF PEER COMPOSITION ON FIRM PERFORMANCE

Panel A: Main performance measures							
Dependent var.:	log Sales (1)	Profit (10,000 RMB) (2)	log Number of employees (3)	log Total assets (4)	log Material cost (5)	log Utility cost (6)	log Productivity (7)
Post* log peer size	0.105*** (0.040)	27.825** (13.432)	0.043 (0.032)	-0.016 (0.034)	0.100* (0.052)	0.141*** (0.042)	0.029 (0.020)
Post* firm demographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,183	4,076	4,183	4,183	4,148	4,086	4,183
Panel B: Intermediate outcomes and alternative explanations							
Dependent var.:	log Number of clients (8)	log Number of suppliers (9)	Bank loan (10)	Management (11)	Innovation (12)	log Reported - log book sales (13)	Tax/sales (14)
Post* log peer size	0.068** (0.032)	-0.001 (0.030)	0.017 (0.016)	0.162*** (0.027)	0.027 (0.017)	0.022 (0.014)	-0.001 (0.001)
Post* firm demographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	No	No	Yes	Yes
Observations	4,173	4,170	4,183	2,774	1,409	4,152	4,178

Notes. Table only uses data for treated firms. Specification (11) is based only on the midline and endline surveys; specification (12) is based only on the endline survey; in those two specifications we also included uninteracted firm demographics. Log peer size is the average of log employment of other group members. Firm demographics are size category, sector, subregion, and their interactions. Standard errors clustered at the meeting group level in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

Peer effect estimates are not significant for employment, assets, or productivity, but are significant, with magnitudes comparable to the sales effect, for materials and utility costs.

Panel B reports the peer effect coefficients for intermediate outcomes and for outcomes that proxy alternative explanations. In this panel we find significant peer effects for the number of clients and the management score.²² Reassuringly, the final two columns show insignificant and small effects for the difference between reported and book sales and the tax to sales ratio. Altogether, the table presents significant peer effects for 6 of the 12 performance measures we had considered in [Section III.A](#) (not counting the last two outcomes, where we expect zeros). In our view these results constitute strong evidence for peer effects.

1. Specification Checks. We now turn to specification checks for the above estimates. These checks make explicit how we exploit the randomness in group assignment for identification, and address an “exclusion bias” that can invalidate the exclusion restriction in [equation \(4\)](#) even when groups are randomly assigned ([Guryan, Kroft, and Notowidigdo 2009](#); [Caeyers and Fafchamps 2016](#)). The exclusion bias results because, even with random assignment, a firm’s baseline characteristics are slightly negatively correlated with its peers’ baseline characteristics since the firm is left out when we compute the peer characteristic.

We present the results of two specification tests in [Appendix A.2](#). In the first, we estimate a variant of [equation \(4\)](#) in which we use only the surprise component of peer size, which is entirely due to the randomization in group assignment. This specification is explicit in exploiting only the exogenous randomness in peers. Formally, we define surprise peer size as the difference between (log) peer size and its expectation over all possible realizations of the group assignment randomization.²³ Since surprise peer size is by design fully orthogonal to all baseline firm characteristics, using it also addresses the exclusion bias. In [Appendix Table A.3](#) we report peer effect regressions in which we include the interaction of $Post_{it}$ with expected and surprise peer size. The coefficients of

22. For the management score our regression only includes the midline and endline data, and for innovation only the endline data, hence in these specifications we omit firm fixed effects and instead control for our firm demographics.

23. To compute expected peer size, we redraw our group assignment randomization 1,000 times and average peer size across these hypothetical draws. Randomization checks (not reported) confirm that the resulting surprise peer size is uncorrelated with a wide range of baseline firm characteristics.

surprise peer size are very similar to those reported in [Table VII](#), showing significant peer effects (at the 10% level) for 8 of the 12 firm performance measures. These results confirm that our controls in [equation \(4\)](#) succeed in isolating the random component of group assignment, and show that the exclusion bias—which would drive a wedge between the results in [Table VII](#) and [Appendix Table A.3](#)—is small in our setting.

As a second specification test, we estimate a “placebo” regression analogous to [equation \(4\)](#) for control firms, using artificial groups created by a similar procedure to that used to create groups in the treatment. Because meetings were not held by these groups of control firms, we expect no peer effects, but any exclusion bias would still be active. [Appendix Table A.4](#) shows the results: the coefficients are insignificant and small in all specifications, further validating our main specification and further confirming that the exclusion bias is not a major factor.

Taken together, our findings show that peers’ identity matters: randomly assigned better peers generate faster firm growth in several domains. Beyond providing internal validity to our previous estimates, these results also contribute to the large literature on peer effects by establishing such effects in a new domain, managerial interactions, and showing that they influence several firm-level outcomes.²⁴

III.C. Attributing Treatment Effects

Our estimates indicate that the meetings had a large effect on firm performance. Here we discuss a set of potential alternative explanations. While each of these explanations may have contributed to a subset of our results, for each we present (i) evidence indicating that it is unlikely to have been an important factor, and (ii) other evidence that it cannot easily explain. In our view, these facts strongly favor the interpretation that the treatment effects are largely due to performance improvements generated by the meetings.

24. The work on peer effects includes studies about education ([Sacerdote 2001](#); [Carrell, Sacerdote, and West 2013](#)), worker productivity ([Mas and Moretti 2009](#); [Bandiera, Barankay, and Rasul 2010](#)), loan repayment ([Breza 2016](#)), program participation ([Dahl, Loken, and Mogstad 2014](#)), as well as neighborhood effects ([Chetty, Hendren, and Katz 2016](#)), among others. See [Jackson \(2011\)](#) for a review.

1. *Experimenter Demand Effects.* A natural concern is that managers who participated in the meetings felt that they were expected to perform well, and as a result overreported their performance in the midline and endline surveys. Several facts suggest that demand effects are unlikely to drive our results. (i) [Table IV](#) shows that the difference between the self-reported and the book value of sales does not vary with the treatment. It is unlikely that managers would have manipulated the sales number in the book—shown to us, without advance notice that we would ask for it, by the firm’s accountant—because of experimenter demand effects. (ii) [Table III](#) shows significant treatment effects on utility costs, which are not an obvious performance measure and as a result are less likely to be manipulated. (iii) Demand effects are unlikely to have driven the results on peer effects which are identified from variation within the meetings treatment. Those results constitute strong evidence that the meetings had direct economic impact. (iv) Treatment effects persisted one year after the meetings had concluded, while experimenter demand effects should weaken over time.

2. *Improved Access to the Government.* A broad concern is that the meetings improved firm growth not because of interactions between managers, but because of a side effect. One such side effect is that firms in the meetings may have had better access to the government through CIIT. Because—except for the first meeting—managers met without interference from CIIT, there is no obvious forum for regular access to CIIT officials. Since CIIT staff members introduced us to both the treatment and the control firms, it is not clear that treatment firms had better government access than did control firms. Thus the circumstances of the design make this effect unlikely. Improved government access also cannot easily explain the positive peer effects: larger peer firms might have actually crowded out the manager from accessing government officials. We also report in [Online Appendix Table O4](#) peer effect specifications, which show that conditional on peer size, peer government experience was not associated with higher firm growth. Finally, it is not fully clear how access to the government would generate gains in management and innovation.

A second possible side effect is that firms in the meetings could use either the government certificate or the fact of the meetings to signal their quality. This logic cannot work with the formal certificate because it was also given to control firms, and,

as [Table II](#) Panel E shows, there was no difference in managers' willingness to pay for the certificate between treatment and control firms. In addition, because the certificate was only given after the midline survey and there was no obvious way to use it in advance, it is unlikely to have affected the midline results. The signaling logic also cannot explain the positive peer effects.

3. Direct Effect of Government Funding. A variant of the side effect argument is that the effect of the meetings was partly driven by the additional intervention of distributing information about a government funding opportunity. According to this logic, while the meetings helped by facilitating its diffusion, the grant itself generated the performance gains.

Because the grant was decided on and awarded after the midline survey, it could not have directly affected performance at midline.²⁵ Nevertheless there could be effects from simply applying due to the anticipation of winning even at midline, and the grant could have directly affected outcomes at endline.

To explore these issues, in [Table VIII](#) we report estimates that extend our main regression (1) by adding the interactions between the midline and endline indicators and the firm having access to information about the grant. Here access to information is defined to be one for a firm if some member of its meetings group (for treatment firms) or the firm itself (for control firms) was exogenously given information about the funding opportunity. Because the information was randomly provided, this regression is identified. The estimated treatment effects are similar to those in [Table III](#) and show that it is not information about the government grant which drives our main results. The point estimates of the new interactions are small, which is intuitive given that most firms informed about the grant did not win it.²⁶

25. Relatively few firms in our sample won: out of 458 applicants among treatment firms, 37 received funding, whereas out of 218 applicants among control firms, 14 received funding.

26. [Online Appendix](#) Table O5 shows that treatment effect estimates are similar when we directly control for the firm winning the grant. These results are not causal because winning is endogenous, but they probably underestimate the treatment effect because better-performing firms were more likely to win the grant. As another robustness check in [Online Appendix](#) Table O6 we show that information about the grant had no effect on performance in the sample of control firms.

TABLE VIII
EFFECT OF MEETINGS: CONTROLLING FOR INFORMATION ON GOVERNMENT GRANT

Variables	log Sales (1)	Profit (10,000 RMB) (2)	log Number of employees (3)	log Total assets (4)	log Utility cost. (5)	log Number of clients (6)	Management (7)
Midline	-0.010 (0.026)	18.183** (7.825)	0.025 (0.021)	-0.012 (0.021)	-0.045 (0.027)	0.009 (0.023)	
Endline	0.004 (0.038)	12.106 (11.202)	0.038 (0.032)	0.013 (0.038)	0.020 (0.035)	0.041 (0.035)	0.015 (0.039)
Meetings*midline	0.067* (0.037)	30.542** (13.306)	0.058** (0.027)	0.042 (0.032)	0.081** (0.038)	0.085*** (0.031)	0.216*** (0.055)
Meetings*endline	0.091* (0.052)	32.518* (19.087)	0.084* (0.048)	0.099** (0.049)	0.113** (0.048)	0.116** (0.047)	0.240*** (0.052)
Info on funding* midline	0.036 (0.038)	-15.585 (13.741)	-0.017 (0.028)	0.063* (0.035)	0.058 (0.039)	0.017 (0.031)	-0.018 (0.056)
Info on funding* endline	0.023 (0.054)	0.185 (19.21)	-0.022 (0.050)	0.015 (0.051)	0.011 (0.050)	0.009 (0.048)	-0.079 (0.053)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	No
Observations	7,857	7,664	7,857	7,857	7,676	7,841	5,211

Notes: Standard errors are clustered at the meeting group level for treated firms and at the firm level for control firms. Column (7) is based on the midline and endline surveys only. ***p < .01, **p < .05, *p < .1.

4. *Collusion and Business Stealing.* Another potential side effect of the meeting is collusion: perhaps firms in the meetings improved outcomes not because of performance gains but by coordinating price increases. But these firms were small actors in a large market. Also, as emphasized by [Duso, Roller, and Seldeslachts \(2014\)](#), standard models of collusion predict that the increase in profit is accompanied by a reduction in quantity, contradicting the positive treatment effect on factors and inputs. In addition, collusion cannot easily explain other gains, such as improved management.

A variant of this concern is that the impacts were due to business shifting—treatment firms trading with each other at the expense of outsiders—and do not represent aggregate gains. But the results in [Section IV.B](#) below indicate that only about a quarter of the increase in the number of suppliers and clients was due to in-group partnerships. In addition, for the argument to work there has to be a benefit for the firms that shift their business. If this is an economic benefit, then business shifting is just the process of better firms gaining market share through the logic of creative destruction, and should represent aggregate gains. An alternative potential benefit emphasized by [Haselmann, Schoenherr, and Vig \(forthcoming\)](#) is rent extraction. But rents are probably more common in contexts with state-owned firms which lack a direct profit motive than in our context with profit-maximizing private firms. Indeed, much of the crony lending documented by [Haselmann, Schoenherr, and Vig \(forthcoming\)](#) was driven by state-owned banks. In addition, this argument does not explain the gains in management or innovation. Overall, we think that inefficient business shifting was not a major factor in our context.

Based on this discussion, we believe that the most plausible alternative explanations are unlikely to drive our results, and we conclude that the meetings treatment indeed significantly improved firm performance.

IV. MECHANISMS

In this section we use the additional interventions to document two mechanisms operating in the meetings: learning and partnering. Importantly, other mechanisms may have also contributed to the effect of the meetings.

IV.A. Learning

We show that the meetings facilitated the diffusion of business-relevant information using the intervention in which we distributed information about two financial products (independently) to randomly chosen managers. The first product was a firm funding opportunity in the form of a government grant. Because it could be used to improve a competitor's business, we expected that managers would view this product to be rival. The second product was a private savings opportunity: a high-yield investment. Because it would only affect a peer's personal finances, we expected that managers would view this product to be less rival.²⁷ As discussed in Section II.B, we randomized the information about the two products independently and provided it to the same share of treatment and control firms.

1. *Empirical Strategy.* We use two main regressions. First, using the full sample of treatment and control firms in the midline, we estimate, separately for each financial product:

$$(5) \quad \begin{aligned} Applied_i = & const + \gamma_1 \cdot Info_i + \gamma_2 \cdot (1 - Info_i) \times Meetings_i \\ & + \gamma_3 \cdot Info_i \times Meetings_i + \varepsilon_i. \end{aligned}$$

Here the dependent variable is an indicator for whether the manager reported in the midline survey to have applied for the product. The coefficient γ_1 measures whether the information treatment "worked" in increasing the likelihood of application. The coefficient γ_2 measures whether uninformed managers in the meetings treatment were more likely to apply than uninformed managers in the control. A positive γ_2 may indicate information diffusion from peers, some of whom were exogenously informed about the product. But it could also indicate higher demand for funding due to the growth effect of the meetings. γ_3 measures whether the effect of information on applications was higher in the meetings treatment: whether the meetings complemented the effect of information, perhaps through encouragement from peers.

27. Both products were in limited supply. For the funding product 676 firms in our sample applied (458 treatment, 218 control) and 51 won (37 treatment, 14 control). For the saving product 1,653 managers in our sample applied (990 treatment, 663 control); we do not have data on the number of managers who got the product.

To get a more precise measure of diffusion, our second regression uses only the sample of uninformed managers in the meetings treatment in the year after the intervention:

$$\begin{aligned} Applied_i = & const + \gamma_4 \cdot Groupmember\ informed_i + \gamma_5 \cdot Competition_i \\ & + \gamma_6 \cdot Groupmember\ informed_i \times Competition_i + controls \\ (6) \quad & + \varepsilon_i. \end{aligned}$$

Here *Groupmember informed_i* is an indicator of *i* having at least one peer in his or her group who had received the information about the product. Given that the information treatment was randomized, γ_4 measures the causal effect of having an informed peer on the decision to apply. *Competition_i* is an indicator for a higher-than-median level of product market competition in the group of *i*. We define this variable by first computing the average number of in-group competitors of firms in a group (self-reported at midline); and then splitting the set of groups by the median of this value.²⁸ Thus γ_5 measures whether average application rates were lower in more competitive groups, and γ_6 the extent to which diffusion was weaker in more competitive groups. The controls are our usual firm demographics: indicators for subregion, sector categories, and size categories at baseline, and their interactions. Because the randomization into groups was conditioned on these variables, by including them we isolate the variation in *Competition_i* which is driven by the random variation in group composition.

2. *Results.* Table IX presents results about the diffusion of the firm funding opportunity. The first two columns show the estimates from regression (5). Column (1), which only includes *Info_i*, shows that being informed increased the likelihood of application by a highly significant 30 percentage points. This confirms that the information treatment worked. Column (2) also includes the interactions with the meetings treatment. Among uninformed managers, being in the meetings treatment increased application rates by a highly significant 20.2 percentage points. This effect can

28. Two facts justify the use of the self-reported designations to identify competitors. First, over 90% of the competitor designations were reciprocated. Second, on average 98% of the peers designated as competitors, but only 35% of the peers not designated as competitors, were in the same two-digit industry as the firm.

TABLE IX
DIFFUSION OF INFORMATION ABOUT FUNDING OPPORTUNITY FOR THE FIRM

Dependent var.:	Applied for the firm funding product				
	(1)	(2)	(3)	(4)	(5)
Sample:	All firms		Uninformed firms in meetings		
Info	0.300*** (0.021)	0.370*** (0.023)			
No info * meetings		0.202*** (0.025)			
Info * meetings		0.072** (0.032)			
Having informed group members			0.291*** (0.035)		0.411*** (0.054)
Competition				-0.150*** (0.052)	-0.060 (0.040)
Having informed group members *competition					-0.212*** (0.068)
Firm demographics	No	No	Yes	Yes	Yes
Observations	2,646	2,646	846	846	846

Notes. Table uses data from the midline survey. Competition is 1 for groups in which the average number of competitors (reported by firms) is higher than the median across groups, and 0 otherwise. Firm demographics are firm size category, sector, subregion, and their interactions. Standard errors clustered at the meeting group level in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

come either from information diffusion or from increased demand for funding because of firm growth. More surprisingly, among informed managers the meetings treatment also increased the probability of application by a significant 7.2 percentage points. Thus in our context formal funding and business networks complemented each other, a result that may be viewed as a positive interaction between formal and informal institutions (Fafchamps 2016).

The remaining columns of the table report estimates of regression (6). The significant coefficient of 0.291 in column (3) shows that having at least one informed group member increased the probability of application by 29.1 percentage points. This result provides direct causal evidence that the meetings diffused information, that is, the learning channel. Column (4) suggests that on average competition reduced application rates. In column (5) the significant and negative interaction effect of -0.212 suggests that competition reduced the strength of information diffusion about the firm funding product. Intuitively, managers may have been less willing to share rival information with their competitors. Overall, these results show that the meetings channeled

TABLE X
 DIFFUSION OF INFORMATION ABOUT SAVING OPPORTUNITY FOR THE MANAGER

Dependent var.:	Applied for the private saving product				
	(1)	(2)	(3)	(4)	(5)
Sample:	All firms		Uninformed firms in meetings		
Info	0.398*** (0.018)	0.542*** (0.023)			
No info * meetings		0.276*** (0.028)			
Info * meetings		0.007 (0.022)			
Having informed group members			0.346*** (0.033)		0.341*** (0.048)
Competition				0.005 (0.046)	0.018 (0.046)
Having informed group members * Competition					0.016 (0.065)
Firm demographics	No	No	Yes	Yes	Yes
Observations	2,646	2,646	835	835	835

Notes. Table uses data from the midline survey. Competition is 1 for groups in which the average number of competitors (reported by firms) is higher than the median across groups, and 0 otherwise. Firm demographics are firm size category, sector, subregion, and their interactions. Standard errors clustered at the meeting group level in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

business-relevant information and also suggest that diffusion was mediated by the extent of competition.²⁹

In Table X we explore the diffusion of information about the private savings opportunity. The structure is identical to that of the previous table. Column (1) shows that the information treatment was very effective for this product as well, and column (2) shows that there was no complementarity between networks and a personal financial product. Column (3) presents direct evidence for information diffusion, while columns (4) and (5) suggest that competition did not affect the strength of diffusion. Consistent with our prior expectation that it is less rival, the estimates suggest that competition did not influence the diffusion of information about this product. The fact that competitive groups had lower diffusion only for the rival product supports the interpretation that lower diffusion was driven by the unwillingness to share rival

29. The fact that we find positive diffusion even in the competitive groups ($0.41 - 0.21 = 0.2 > 0$) suggests—similarly to the model of Stein (2008)—that the benefits of sharing knowledge exceeded the cost of helping competitors in our context.

information with competitors, not some correlated omitted factor that generally reduced information diffusion.³⁰

The interpretation that competition reduces diffusion rates raises the question of why small firms in a large market worry about competitors getting funding. We see two possible answers here. First, a competitor familiar with the firm's operations may be able to use the grant to steal the firm's business by targeting its clients with lower-priced offers. This risk can act as an incentive for the firm's manager not to share information about the grant. Second, for some firms a form of context effect may be active (Kamenica 2008): the manager may generally feel cautious about helping a competitor's business even if the concrete action does not generate a clear and direct business loss to her firm. Note that neither of these arguments contradicts the logic in Section III.C that firms were unlikely to collude, because if they were to jointly raise prices clients could still choose other sellers.

The results also raise the question of whether in more competitive groups the overall gains from the meetings were smaller. To explore this, in Online Appendix Table O8 we report results from a peer effect specification asking whether firms gained more in groups with fewer competitors. We find insignificant effects on all outcomes. Possible explanations include the negative effects of competition being small relative to the benefits from the meetings, or offsetting effects such as similar firms being better sources of advice.

Overall, we interpret the results as direct evidence on the learning-from-peers channel. Beyond highlighting a concrete mechanism of the meetings, our findings also inform a literature studying information diffusion in social networks.³¹ Our contribution to this work is to demonstrate the effect—also explored theoretically in a model by Immorlica, Lucier, and Sadler (2014)—that competition may limit the diffusion of rival information. In

30. Online Appendix Table O7 shows that the diffusion rate was similarly high in the 50% and the 80% information treatment: there was no major increase in uninformed managers' probability of application when having 80% rather than 50% informed group members.

31. Much of this work has explored the diffusion of technology (Bandiera and Rasul 2006; Conley and Udry 2010), and financial choices (Duflo and Saez 2003; Hong, Kubik, and Stein 2004; Banerjee et al. 2013; Cai, de Janvry, and Sadoulet 2015) in the social networks of individuals. More recent work on the diffusion of business choices in managerial networks includes Cohen, Frazzini, and Malloy (2008), who study the diffusion of financial information, and Fafchamps and Quinn (forthcoming), who study the diffusion of certain management practices.

combination with [Hardy and McCasland \(2016\)](#), who show in independent work that the diffusion of a new weaving technique in Ghana was lower in treatments with higher experimentally induced competition, these results highlight a novel friction in technology diffusion: the endogenous (dis)incentive to transmit information. This friction may generate a new, as yet unexplored interaction between technology spillovers and product market rivalry ([Bloom, Schankerman, and Van Reenen 2013](#)).

IV.B. Partnering

We use the cross-group intervention to document evidence on the partnering mechanism. Our approach is to compare the number of new connections in the regular groups and in the cross-groups. This comparison is relevant for two reasons. First, it can attribute some of the increase in partnerships documented earlier ([Table IV](#)) to a reduction in partnering costs created by the regular meetings, that is, the partnering mechanism. Indeed, if the increase in partnerships was only due to other mechanisms, such as treatment-induced firm growth, then we expect no difference in partnering in the regular versus the cross groups. Second, the comparison can reveal whether the friction in partnering was only lack of information about the identity of potential partners, which seems to be a key friction studied in search-and-matching models of the labor market ([Rogerson, Shimer, and Wright 2005](#)). If that was the only friction, then again we expect no difference in partnering rates.

We compare relationships in the regular and the cross-groups using the regression

$$(7) \quad \begin{aligned} Relation_{igt} = & \text{const} + \theta_1 \cdot \text{Midline}_{igt} \times \text{Regular}_{igt} + \theta_2 \cdot \text{Endline}_{igt} \\ & \times \text{Regular}_{igt} + \text{Controls} + \text{Firm f.e.} + \varepsilon_{igt}. \end{aligned}$$

Here each observation is a firm, group category (regular or cross), and year triple. The sample consists of observations in the midline and endline waves for the set of firms that participated in both regular and cross-group meetings. The dependent variable is a measure of relationships between firm i and peers in group g in year t , such as the number of active partners from the group in that year. The coefficients of interest are θ_1 and θ_2 , which measure the extent to which firms had more relationships with peers in the regular group. For controls we include the share of

TABLE XI
REPEATED INTERACTIONS AND NEW PARTNERSHIPS

Variables	Number of referrers (1)	Number of direct partners (2)	Choice in trust game (3)
Regular meetings*midline	2.178*** (0.119)	1.161*** (0.106)	2.742*** (0.172)
Regular meetings*endline	2.400*** (0.122)	1.275*** (0.107)	3.009*** (0.175)
Peer demographics	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes
Observations	1,744	1,744	1,744
Mean dep. var. for cross-group	0.084	0.302	0.960

Notes. Each observation is a (firm, group category, year) triple. The sample consists of treated firms that participated in both regular and cross-group meetings. Referrer is a group member who referred a partner or employee to the firm in the given year. Direct partner is a group member doing business with the firm in the given year. Peer demographics are the share of peers in the given group which are larger than the subregion median (measured with employment at baseline) and the share of peers in the given group that are in the same sector as the firm. Standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

same-sector firms and the share of firms with size above the subregion median to pick up any variation in group composition which may also affect relationships.³² Table XI reports the results. Column (1) focuses on the number of referrers—peers who referred suppliers, clients, partners, and lower-ranking managers. At midline, on average each manager had a significant 2.18 more peers act as referrers in the regular group than in the cross group. At endline—that is, only counting referrals taking place in the year after the midline survey—the corresponding difference was 2.4. Thus peers in the regular group provided more referrals and continued to do so after the conclusion of the meetings. Column (2) shows the result for the number of direct business partners: suppliers, clients, and firms engaging in other joint business activities such as joint projects. The average manager had a significant 1.16 more active partnerships from the regular group than from the cross group during the year of the program, and a significant 1.28 more active partnerships during the year after the program. Column (3) reports average giving in hypothetical trust games played with a randomly chosen member of the regular group and of the cross group. Managers exhibited significantly more trusting

32. Balance checks (not reported) show that group composition measures including the baseline average size and sector of peer firms were not significantly different between the regular and cross groups.

behavior towards their peers in the regular group at midline and endline.³³

Our results indicate that the meetings reduced the cost of referrals and partnerships, so that partnering is indeed one of the channels through which they improved firm performance. Moreover, referrals and partnerships continued to be active in the year after the conclusion of the meetings. The result on trust game behavior suggests that these lower partnering costs may have emerged in part because repeated meetings created trust between managers. We conclude that lack of trust is likely to be an important barrier to creating business partnerships in our context.

These results contribute to a literature that studies network-based referrals in the labor market by documenting referrals in a new domain: managers referring business partners.³⁴ Our result on trust relates to the research about trust in networks. [Karlan et al. \(2009\)](#) show theoretically that networks that embed more trust are more useful for making high-value referrals, while [Feigenberg, Field, and Pande \(2013\)](#) establish that regular meetings between microfinance borrowers build trust and improve loan performance. Our findings are consistent with these results and highlight the importance of trust in firm-to-firm interactions.

Taken together, our results on learning and partnering suggest that the meetings created some of the benefits which are commonly associated with business clusters ([Porter 1998](#)), but without the firms actually moving near each other.³⁵

V. CONCLUSION

In this article we used a field experiment with experimental business associations to measure the effect of business networks

33. We used the following trust question. "Suppose that you are given RMB 100,000. Out of this, you can choose to give as much as you want for a business project which is controlled by person X. This project is very successful and triples the money you give. All the proceeds go to person X. Person X can then choose to return to you as much of the money the project earns as he wishes. How much (between 0 and RMB 100,000) do you give to person X?"

34. [Calvo-Armengol and Jackson \(2004\)](#) is a model of network-based job referrals, while [Ioannides and Loury \(2004\)](#) document evidence on their role in the labor market.

35. Recent work on production and entrepreneurial clusters includes [Guiso and Schivardi \(2007\)](#), [Martin, Mayer, and Mayneris \(2011, 2013\)](#) and [Guiso, Pistaferri, and Schivardi \(2015\)](#).

on firm performance. We found significant, robust, large, and persistent effects of the meetings on sales, profits, factors, inputs, innovation, and management. We also found direct evidence on two mechanisms, learning and partnering. We now discuss some implications of these results.

We begin with a cost benefit analysis. Combining publicly available survey and wage growth data, we estimate the annual wages of managers in our sample to be RMB 812,300.³⁶ This value accords with the range locals reported to us informally. We assume that all reported profits also accrue to the manager, for an additional RMB 800,000 on average. If managers work 200 days a year and each meeting takes a full day, the time cost of the meetings is about RMB 98,000 for our average manager. Additional costs include the certificate, the cost of which we assume is at most twice its value to the average manager (less than $2 \times$ RMB 6,000), the cost of the government funding opportunity per manager in the treatment group (RMB 5,000), and the organizational costs of recruiting and arranging the meetings, estimated by CIIT to be RMB 2,500.³⁷ The total estimated cost per manager is thus RMB 117,000. As Table III shows, the average annual profit gain by the midline survey was about RMB 250,000, more than twice the estimated cost. Although there is clearly noise in these calculations, the result strongly suggests that the meetings were quite cost-effective.³⁸

Given this result, a natural question is why the managers did not organize meetings for themselves. There are several possible answers. First, search costs and trust barriers may be higher if managers were to organize the meetings themselves: they would need to find—without the help of CIIT—others willing to form groups with unfamiliar people. Second, there may be a public good problem if these costs of organization fall on a single manager.

36. A survey conducted by the All-China Federation of Industry and Commerce shows the average earning of private business owners to be about RMB 200,000 in 2005. We multiplied this value with wage growth in the private sector between 2005 and 2014 (a factor of 4.06 by the Chinese National Bureau of Statistics) to obtain our estimate. A summary report of the survey is available at <http://www.people.com.cn/GB/jingji/42775/3164559.html>.

37. We estimate the cost of government funding by rounding up the product of the maximum amount (RMB 200,000) and the share of treatment firms receiving it (37 out of 1,500).

38. We do not include researcher time or the survey costs in the calculation because they are not required to implement the design.

Third, paralleling the argument of [Bloom et al. \(2013\)](#), managers may have underestimated the gains from business associations or from changing business practices. Each of these explanations suggests that managers should continue to interact with their newly found peers after the conclusion of our intervention. This is indeed what the endline survey shows: during the year after the conclusion of the meetings, 57% of treated managers reported to have met at least one group member on average once a month, and 81% reported to have met at least one group member on average once every two months. We also note that similar business associations—such as the Lions Club or the Rotary Club—exist in more advanced countries, suggesting that at a higher level of economic development the market can sometimes overcome the matching frictions.

We next compare our results to the impacts found in other types of interventions. [McKenzie and Woodruff \(2014\)](#) review several studies evaluating business training and business consulting interventions. For business training they conclude that—perhaps because of limited power—most studies do not find a significant impact on sales or profits.³⁹ In contrast, the high-intensity management consulting intervention evaluated by [Bloom et al. \(2013\)](#) did create a large productivity increase of 17%. Our 8% sales effect is smaller than this; but our intervention is cheaper, easier to implement, and quite cost-effective. Finally, [Brooks, Donovan, and Johnson \(forthcoming\)](#) show that a one month business mentoring intervention for Kenyan microenterprises led to a 20% profit effect, which faded the year after the intervention. The mechanisms they emphasize are similar to the ones we document, but our effects persisted after one year. We conclude that our business meetings intervention had surprisingly large effects in comparison to other interventions that have been evaluated.

Which aspects of the design made the intervention successful? The comparison with the designs of other studies, and the results on mechanisms, allow us to formulate some hypotheses. First, similar to the [Bloom et al. \(2013\)](#) study, but unlike many business

39. Exceptions include [Calderon, Cunha, and de Giorgi \(2013\)](#), who find a 20% impact on sales and a 24% impact on profits, and [De Mel, McKenzie, and Woodruff \(2014\)](#), who find a 41% increase in sales and a 43% increase in profits for start-up businesses. But these estimates are quite noisy. Our sales and profit impacts fall within their standard error bands, are more precisely estimated, and are persistent.

training evaluations, our sample of firms was selected. This fact suggests that firm interventions are more likely to succeed when managers themselves are interested in improving their business and that a possible way to identify such “gazelles” (Fafchamps and Woodruff 2017) may be to use an explicit recruitment process.⁴⁰ Second, also paralleling the Bloom et al. (2013) study, our intervention was quite intensive. Managers met every month and combined company visits with hours of discussion. This intensity may have contributed in multiple ways. The results on meeting frequency suggest that it helped build trust. Observing other firms’ operations in practice may have enhanced learning through a demonstration effect. Third, our results on management and information diffusion suggest that managers had gaps of knowledge that learning could fill. This could be because the firms were young and did not have access to other sources of business information.

This discussion suggests that the following conditions on the design increase the probability of a successful business meetings policy. (i) Self-selected pool of firms. (ii) Regular intensive meetings involving firm visits. (iii) Young firm age. The discussion also suggests that meetings are more likely to help in contexts in which the following distortions are important. (iv) Contracting problems, which increase the value of trust. (v) Relative lack of alternative sources of business information (e.g., MBA programs). We hope that these conditions can help guide future interventions and scale-ups and thus contribute to private sector development.

APPENDIX

A.1. Sample Selection Checks

We present two tables relevant for the discussion in Section II.D about sample selection. Appendix Table A.1 addresses selective attrition by showing that the baseline characteristics of firms that attrited—either at midline or at endline—were not significantly different by treatment status. Appendix Table A.2 shows that at baseline nonapplicant firms were generally smaller than applicant firms.

40. In the context of our meetings program, recruiting good firms has the direct benefit that they may respond to the treatment and the indirect benefit that—through peer effects—they generate higher growth for other participants.

APPENDIX TABLE A.1
 BASELINE CHARACTERISTICS OF ATTRITING FIRMS

	Treatment	Control	Difference
Number of observations	139	117	
Firm age	2.39 (1.82)	2.58 (1.58)	-0.19 (0.22)
Sector: manufacturing	0.50 (0.50)	0.50 (0.50)	-0.00 (0.06)
Number of employees	30.43 (48.89)	38.23 (60.69)	-7.80 (6.85)
Bank loan (1=yes, 0=no)	0.25 (0.43)	0.27 (0.44)	-0.02 (0.06)
Log sales	5.44 (3.17)	5.79 (3.06)	-0.36 (0.39)
Gender (1=male, 0=female)	0.84 (0.37)	0.89 (0.32)	-0.05 (0.04)
Age	37.46 (16.24)	38.49 (15.07)	-1.03 (1.97)
Education: college	0.29 (0.45)	0.30 (0.46)	-0.01 (0.06)
Government working experience	0.22 (0.41)	0.27 (0.45)	-0.06 (0.05)

Notes. Table shows baseline summary statistics for firms that ever attrited. Standard deviations are in parentheses for the Treatment and Control columns. The difference column reports the difference in characteristics between the treatment and control groups, and standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

A.2. Peer Effect Specification Tests

We report two specification tests for the main peer effect results in [Section III.B](#). Our first regression is

$$(8) \quad y_{it} = \text{const} + \delta_1 \cdot \text{Post}_{it} + \delta_2 \cdot \text{Post}_{it} \times \text{Surprise log Peer size}_{it} \\ + \delta_3 \cdot \text{Post}_{it} \times \text{Expected log Peer size}_{it} + \text{Firm f.e.} + \varepsilon_{it}.$$

Here expected log peer size (at baseline) is the expectation taken over all possible realizations of the group assignment randomization, and surprise log peer size is the difference between log peer size and its expectation. Thus δ_2 measures the effect of the purely random component of peer size—explicitly created by the randomization of the intervention—on firm performance. [Appendix Table A.3](#) reports the results, and shows significant effects in eight

APPENDIX TABLE A.2.
COMPARING SAMPLE FIRMS WITH NONAPPLICANT FIRMS

	Sample firms	Nonapplicant firms	Difference
Number of observations	2,820	124	
Panel A: Firm characteristics			
Number of employees	36.19 (86.49)	18.43 (21.44)	17.76** (7.78)
Sales (10,000 RMB)	1,593.70 (6,475.18)	548.30 (705.87)	1,044.40* (581.75)
Net profit (10,000 RMB)	79.23 (205.35)	25.84 (36.16)	53.39*** (18.46)
Bank loan (1=yes, 0=no)	0.25 (0.43)	0.21 (0.41)	0.04 (0.04)
Panel B: Managerial characteristics			
Gender (1=male, 0=female)	0.84 (0.37)	0.73 (0.44)	0.11*** (0.03)
Age	40.84 (8.85)	42.51 (9.18)	-1.67*** (0.81)
Communist Party member (1=yes, 0=no)	0.21 (0.40)	0.14 (0.34)	0.07* (0.04)

Notes. The first two columns show baseline characteristics of firms in our sample, and of 124 randomly chosen nonapplicant firms. Standard deviations are in parentheses. The third column reports the difference in characteristics, and standard errors in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

of the twelve firm performance outcomes (specifications (1)–(12)) as well as zero effects in the two placebo outcomes (specifications (13)–(14)). These results are very similar to those in [Table VII](#), validating our main specification and showing that the exclusion bias in our setting is small.

[Appendix Table A.4](#) reports results from estimating [equation \(4\)](#) among control firms, using artificial groups created by a procedure similar to that used to create the groups in the treatment. All coefficients are insignificant, further validating our main specification and further supporting that the exclusion bias is small in our setting.

APPENDIX TABLE A.3.
PEER COMPOSITION EFFECT: USING THE SURPRISE COMPONENT OF PEER SIZE

Panel A: Main performance measures								
Dependent var.:	log Sales	Profit (10,000 RMB)	log Number of employees	log Total assets	log Material cost	log Utility cost	log Productivity	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	
Post* surprise log peer size	0.095** (0.044)	37.961** (14.618)	0.013 (0.032)	-0.029 (0.034)	0.068* (0.052)	0.127*** (0.045)	0.037* (0.021)	
Post* expected peer size	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Observations	4,183	4,076	4,183	4,183	4,148	4,086	4,183	
Panel B: Intermediate outcomes and alternative explanations								
Dependent var.:	log Number of clients	log Number of suppliers	Bank loan	Management	Innovation	log Reported - log book sales	Tax/sales	
	(8)	(9)	(10)	(11)	(12)	(13)	(14)	
Post* surprise log peer size	0.077** (0.036)	0.002 (0.034)	0.017 (0.015)	0.166*** (0.027)	0.028 (0.017)	0.020 (0.014)	-0.001 (0.001)	
Post* expected peer size	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Firm fixed effects	Yes	Yes	Yes	No	No	Yes	Yes	
Observations	4,173	4,170	4,183	2,774	1,409	4,152	4,178	

Notes. Table only uses data for treated firms. Specification (11) is based only on the midline and endline surveys; specification (12) is based only on the endline survey; in those two specifications we also included uninteracted firm demographics. Surprise log peer size is the difference between log peer size and its expectation, the latter computed as the average across all realizations of the group assignment randomization. Standard errors clustered at the meeting group level are in parentheses. ***, ***, *p < .01, **p < .05, *p < .1.

APPENDIX TABLE A.4.
PLACEBO EFFECT OF PEER COMPOSITION: CONTROL FIRMS

Panel A: Main performance measures									
Dependent var.:	log Sales (1)	Profit (10,000 RMB) (2)	log Number of employees (3)	log Total assets (4)	log Material cost (5)	log Utility cost (6)	log Productivity (7)		
Post* log peer size	0.031 (0.026)	13.737 (9.250)	-0.022 (0.023)	-0.017 (0.027)	0.066 (0.093)	0.027 (0.025)	0.002 (0.014)		
Post* firm demographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes		
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes		
Observations	3,671	3,586	3,671	3,671	3,641	3,587	3,671		
Panel B: Intermediate outcomes and alternative explanations									
	log Number of clients (8)	log Number of suppliers (9)	Bank loan (10)	Management (11)	Innovation (12)	log Reported - log book sales (13)	Tax/sales (14)		
Post* log peer size	0.022 (0.029)	-0.034 (0.031)	-0.010 (0.015)	0.012 (0.031)	0.010 (0.018)	-0.007 (0.008)	-0.001 (0.001)		
Post* firm demographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes		
Firm fixed effects	Yes	Yes	Yes	No	No	Yes	Yes		
Observations	3,665	3,653	3,671	2,435	1,236	3,642	3,668		

Notes. Table only uses data for control firms. Groups are artificial groups that were created by a similar procedure to the one used in the treatment but did not meet. Specification (1) is based only on the midline and endline surveys; specification (12) is based only on the endline survey; in those two specifications we also included uninteracted firm demographics. Log peer size is the average of log employment of other group members. Firm demographics are size category, sector, subregion, and their interactions. Standard errors clustered at the meeting group level in parentheses. *** $p < .01$, ** $p < .05$, * $p < .1$.

UNIVERSITY OF MARYLAND, NATIONAL BUREAU OF ECONOMIC RESEARCH, AND BUREAU FOR RESEARCH AND ECONOMIC ANALYSIS OF DEVELOPMENT
CENTRAL EUROPEAN UNIVERSITY AND CENTER FOR ECONOMIC AND POLICY RESEARCH

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online. Data and code used to generate tables and figures in this article can be found in Cai and Szeidl (2017), in the Harvard Dataverse, doi:10.7910/DVN/5ZX8ZI.

REFERENCES

- Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80 (2012), 1977–2016.
- Antràs, Pol, and Davin Chor, "Organizing the Global Value Chain," *Econometrica*, 81 (2013), 2127–2204.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul, "Social Incentives in the Workplace," *Review of Economic Studies*, 77 (2010), 417–459.
- Bandiera, Oriana, and Imran Rasul, "Social Networks and Technology Adoption in Northern Mozambique," *Economic Journal*, 116 (2006), 869–902.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson, "The Diffusion of Microfinance," *Science*, 341 (2013), 363–371.
- Banerjee, Abhijit, and Esther Duflo, "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program," *Review of Economic Studies*, 81 (2014), 572–607.
- Bernard, Andrew B., Andreas Moxnes, and Yukiko U. Saito, "Production Networks, Geography and Firm Performance," *Journal of Political Economy*, forthcoming.
- Bernstein, Shai, Xavier Giroud, and Richard R. Townsend, "The Impact of Venture Capital Monitoring," *Journal of Finance*, 71 (2016), 1591–1622.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 128 (2013), 1–51.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81 (2013), 1347–1393.
- Bloom, Nicholas, and John Van Reenen, "Measuring and Explaining Management Practices across Firms and Countries," *Quarterly Journal of Economics*, 122 (2007), 1351–1408.
- Breza, Emily, "Peer Effects and Loan Repayment: Evidence from the Krishna Default Crisis," Working paper, Columbia University, 2016.
- Brooks, Wyatt, Kevin Donovan, and Terence R. Johnson, "Mentors or Teachers? Microenterprise Training in Kenya," *American Economic Journal: Applied Economics*, forthcoming.
- Bruhn, Miriam, Dean Karlan, and Antoinette Schoar, "The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico," *Journal of Political Economy*, forthcoming.
- Caeyers, Bet, and Marcel Fafchamps, "Exclusion Bias in the Estimation of Peer Effects," Working paper, Institute for Fiscal Studies and Stanford University, 2016.

- Cai, Jing, Alain de Janvry, and Elisabeth Sadoulet, "Social Networks and the Decision to Insure," *American Economic Journal: Applied Economics*, 7 (2015), 81–108.
- Cai, Jing, and Adam Szeidl, "Replication Data for: 'Interfirm Relationships and Business Performance'," 2017, *Harvard Dataverse*, doi:10.7910/DVN/5ZX8ZI.
- Calderon, Gabriela, Jesse Cunha, and Giacomo de Giorgi, "Business literacy and development: Evidence from a Randomized Trial in Rural Mexico," NBER Working Paper no. 19740, 2013.
- Calvo-Armengol, Antoni, and Matthew O. Jackson, "The Effects of Social Networks on Employment and Inequality," *American Economic Review*, 94 (2004), 426–454.
- Carrell, Scott, Bruce Sacerdote, and James E. West, "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation," *Econometrica*, 81 (2013), 855–882.
- Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review*, 106 (2016), 855–902.
- Cohen, Lauren, Andrea Frazzini, and Christopher Malloy, "The Small World of Investing: Board Connections and Mutual Fund Returns," *Journal of Political Economy*, 116 (2008), 951–979.
- Conley, Timothy G., and Christopher R. Udry, "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, 100 (2010), 35–69.
- Dahl, Gordon B., Katrine V. Loken, and Magne Mogstad, "Peer Effects in Program Participation," *American Economic Review*, 104 (2014), 2049–2074.
- De Loecker, Jan, "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity," *Econometrica*, 79 (2011), 1407–1451.
- De Mel, Suresh, David McKenzie, and Christopher Woodruff, "Returns to Capital in Microenterprises: Evidence from a Field Experiment," *Quarterly Journal of Economics*, 123 (2008), 1329–1372.
- , "Business Training and Female Enterprise Start-up, Growth, and Dynamics: Experimental Evidence from Sri Lanka," *Journal of Development Economics*, 106 (2014), 199–210.
- Duflo, Esther, and Emmanuel Saez, "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," *Quarterly Journal of Economics*, 118 (2003).
- Duranton, Gilles, and Diego Puga, "Micro-Foundations of Urban Agglomeration Economies," in *Handbook of Regional and Urban Economics*, vol. 4: *Cities and Geography*, J. Vernon Henderson and Jacques-Francois Thisse, eds. (Amsterdam: Elsevier, 2004). 2063–2117.
- Duso, Tomaso, Lars-Hendrik Roller, and Jo Seldeslachts, "Collusion through Joint R&D: An Empirical Assessment," *Review of Economics and Statistics*, 96 (2014), 349–370.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz, "Firm-to-Firm Trade: Imports, Exports, and the Labor Market," Working paper, 2015.
- Fafchamps, Marcel, "Formal and Informal Market Institutions: Embeddedness Revisited," Working paper, Stanford University, 2016.
- Fafchamps, Marcel, and Simon Quinn, "Networks and Manufacturing Firms in Africa: Results from a Randomized Field Experiment," *World Bank Economic Review*, forthcoming.
- Fafchamps, Marcel, and Christopher Woodruff, "Identifying Gazelles: Expert Panels vs. Surveys as a Means to Identify Firms with Rapid Growth Potential," *World Bank Economic Review*, 31 (2017), 670–686.
- Feigenberg, Benjamin, Erica Field, and Rohini Pande, "The Economic Returns to Social Interaction: Experimental Evidence from Microfinance," *Review of Economic Studies*, 80 (2013), 1459–1483.
- Giorcelli, Michela, "The Long-Term Effects of Management and Technology Transfers," Working paper, Stanford University, 2017.

- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi, "Learning Entrepreneurship from Other Entrepreneurs?," NBER Working Paper no. 21775, 2015.
- Guiso, Luigi, and Fabiano Schivardi, "Spillovers in Industrial Districts," *Economic Journal*, 117 (2007), 68–93.
- Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo, "Peer Effects in the Workplace: Evidence from Random Groupings in Professional Golf Tournaments," *American Economic Journal: Applied Economics*, 1 (2009), 34–68.
- Hardy, Morgan, and Jamie McCasland, "It Takes Two: Experimental Evidence on the Determinants of Technology Diffusion," Working paper, NYU Abu Dhabi and University of British Columbia, 2016.
- Haselmann, Rainer, David Schoenherr, and Vikrant Vig, "Rent Seeking in Elite Networks," *Journal of Political Economy*, forthcoming.
- Hochberg, Yael V., Alexander Ljungqvist, and Yang Lu, "Whom You Know Matters: Venture Capital Networks and Investment Performance," *Journal of Finance*, 62 (2007), 251–301.
- Hong, Harrison, Jeffrey D. Kubik, and Jeremy C. Stein, "Social Interaction and Stock-Market Participation," *Journal of Finance*, 59 (2004), 137–163.
- Hsieh, Chang-Tai, and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124 (2009), 1403–1448.
- Immorlica, Nicole, Brendan Lucier, and Evan Sadler, "Sharing Rival Information," Working paper, New York University, 2014.
- Ioannides, Yannis M., and Linda Datcher Loury, "Job Information Networks, Neighborhood Effects, and Inequality," *Journal of Economic Literature*, 42 (2004), 1056–1093.
- Jackson, Matthew O., "An Overview of Social Networks and Economic Applications," *Handbook of Social Economics* (2011), 511–585.
- Kamenica, Emir, "Contextual Inference in Markets: On the Informational Content of Product Lines," *American Economic Review*, 98 (2008), 2127–49.
- Karlan, Dean, Markus Mobius, Tanya Rosenblat, and Adam Szeidl, "Trust and Social Collateral," *Quarterly Journal of Economics*, 124 (2009), 1307–1361.
- Khwaja, Asim Ijaz, Atif Mian, and Abid Qamar, "Bank Credit and Business Networks," HKS Faculty Research Working Paper RWP11-017, 2011.
- Martin, Philippe, Thierry Mayer, and Florian Mayneris, "Public Support to Clusters: A Firm Level Study of French 'Local Productive Systems'," *Regional Science and Urban Economics*, 41 (2011), 108–123.
- , "Are Clusters More Resilient in Crises? Evidence from French Exporters in 2008–2009," DEPR discussion paper, 2013.
- Mas, Alexandre, and Enrico Moretti, "Peers at Work," *American Economic Review*, 99 (2009).
- McKenzie, David, "Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition," *American Economic Review*, forthcoming (2017).
- McKenzie, David, and Christopher Woodruff, "What Are We Learning from Business Training and Entrepreneurship Evaluations around the Developing World?," *World Bank Research Observer*, 29 (2014), 48–82.
- McMillan, John, and Christopher Woodruff, "Interfirm Relationships and Informal Credit in Vietnam," *Quarterly Journal of Economics*, 114 (1999), 1285–1320.
- Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71 (2003), 1695–1725.
- Ministerio de la Produccion del Peru, "Mesas Ejecutivas: A New Tool for Productive Diversification," Technical report, 2016.
- Oberfeld, Ezra, "A Theory of Input-Output Architecture," *Econometrica*, forthcoming.
- Porter, Michael, "Clusters and the New Economics of Competition," *Harvard Business Review*, 76 (1998), 77–91.
- Rogerson, Richard, Robert Shimer, and Randall Wright, "Search Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 43 (2005), 959–988.

- Rosenthal, Stuart S., and William C. Strange, "Evidence on the Nature and Sources of Agglomeration Economies," in *Handbook of Regional and Urban Economics*, vol. 4: "Cities and Geography," J. Vernon Henderson and Jacques-Francois Thisse, eds. (Amsterdam: Elsevier, 2004), 2119–2171.
- Sacerdote, Bruce, "Peer Effects with Random Assignment: Results for Dartmouth Roommates," *Quarterly Journal of Economics*, 116 (2001), 681–704.
- Shue, Kelly, "Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers," *Review of Financial Studies*, 26 (2013), 1401–1442.
- Stein, Jeremy, "Conversations among Competitors," *American Economic Review*, 98 (2008), 2150–2162.
- Syverson, Chad, "What Determines Productivity?," *Journal of Economic Literature*, 49 (2011), 326–365.