

Interfirm Relationships and Business Performance*

Jing Cai
University of Michigan and NBER

Adam Szeidl
Central European University and CEPR

September 13, 2016

Abstract

We organize regular business meetings for randomly selected managers of young Chinese firms to study the effect of business networks on firm performance. We randomize 2,800 managers into several groups that hold monthly meetings for one year, and a “no-meetings” control group. We find that: (1) The meetings increase firm revenue by 7.8 percentage points, and also significantly increase profit, a management score, employment, and the number of business partners; (2) These effects persist one year after the conclusion of the meetings; and (3) Firms randomized to have better peers exhibit higher growth. We exploit additional interventions to document concrete channels: (4) Peers share exogenous business-relevant information, particularly when they are not competitors, showing that the meetings facilitate learning; (5) Managers create more business partnerships in the regular than in other one-time meetings, showing that the meetings improve firm-to-firm matching.

*Emails: caijing@umich.edu, szeidla@ceu.edu. We thank Attila Gaspar, Huayu Xu, Hang Yu, and Zhengdong Zhang for excellent research assistance, Daron Acemoglu, Abhijit Banerjee, Andrew Bernard, Nick Bloom, Emily Breza, Arun Chandrasekhar, Esther Dufo, Ben Golub, Matt Jackson, Dean Karlan, David Lam, Ben Olken, Rohini Pande, Mark Rosenweig, Antoinette Schoar, Duncan Thomas, Chris Woodruff, Dean Yang, and seminar participants for helpful comments. We are grateful to Innovations for Poverty Action’s SME Initiative, the Private Enterprise Development in Low-Income Countries, the University of Michigan and the European Research Council for financial support.

1 Introduction

Barriers to firm growth limit economic development. Much research has focused on barriers that operate 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: they rely on business relationships which provide information, training, referrals, intermediate inputs, and many other services. Recent theoretical work has explored how supply chain networks shape industry allocations, and evidence from observational data also supports the importance of business relationships.¹ But we do not fully understand how an exogenous change in business networks affects firm performance, the underlying mechanisms, and the implications for policies that can induce such a change.

To explore these issues, we use a field experiment we conducted in Nanchang, China, in which we organized experimental business associations to managers of young small and medium enterprises (SMEs). We build on the approaches of Fafchamps and Quinn (2014) who generate variation in managerial networks through participation in committees, and Bernard, Moxnes and Saito (2015) who exploit the introduction of the Shinkansen speed train in Japan. In our design, networks are created more directly through regular meetings that have the explicit purpose of fostering business interactions. We also introduce additional interventions to learn about mechanisms. We find that business networks substantially improve firm performance, and show that information diffusion, improved access to partners, trust, and peer effects are all active mechanisms. Because SMEs produce a large share of the output in developing countries, our findings suggest that the policy of organizing business associations can meaningfully contribute to private sector development.

In Section 2 we introduce our experimental design. In the summer of 2013 we invited micro, small and medium enterprises established in the preceding 3 years in Nanchang to participate in business associations. From 2,800 firms which expressed interest, we randomly selected 1,480 and randomized their 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.

¹For example, Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi (2012) study a model of supply chains while McMillan and Woodruff (1999) show that interfirm relationships shape trade credit access in Vietnam. We discuss the literature on firm networks in more detail below.

Managers in each meeting group were encouraged to hold monthly self-organized meetings. These meetings were intensive: managers would typically visit the firm of a group member, spend hours discussing issues they face in running their business, learn from peers, and establish new contacts. The meetings program lasted for one year. As an incentive to participate, after the conclusion of the intervention we gave managers who answered our surveys and attended at least 10 out of the 12 meetings a certificate which provided access to certain government services. We also gave the certificate to managers in the control group who answered our surveys.

We surveyed the firms in 2013 summer before the intervention (baseline), in 2014 summer shortly after the intervention ended (midline), and in 2015 summer one year after the intervention ended (endline). In the surveys we collected information about (1) Firm characteristics including sales, employment, borrowing, and other balance sheet variables; (2) Firm networks and the type of interaction; (3) Managerial characteristics including overconfidence, stress levels, and—in the midline and endline surveys—management practices.

We also introduced three additional interventions to learn about mechanisms. First, to learn about peer effects we created variation in the composition of groups by industry and size. Most of our firms are in two broad industry categories, manufacturing and services, and we categorized firms by their number of employees as “small” or “large”. We then created four kinds of groups: homogenous size and industry; homogenous size and mixed industry; mixed size and homogeneous industry; mixed size and mixed industry.

Second, paralleling the approaches of Duflo and Saez (2003) and Cai, de Janvry and Sadoulet (2015) we provided randomly chosen managers with information about two financial products: (i) a funding opportunity for the firm; (ii) a savings opportunity for the manager. For each product, we created random variation across groups in the share of informed managers. We also provided the information to random control firms to ensure that the same share of treatment and control firms are directly informed.

Third, to learn about the importance of repeated interactions, we organized one-time cross-group meetings for a random subset of managers in the meetings. The cross-groups had 10 managers each, never contained two managers from the same meeting group, and met only once in the spring

of 2014. Moreover, in the midline survey we asked managers to play hypothetical trust games with random members of both their regular- and their cross-group. This intervention is similar to that of Feigenberg, Field and Pande (2013) who studied the effect of meeting frequency for loan performance in microfinance.

In Section 3 we present our results. We first explore the overall impact of the meetings intervention. Our basic regression is a firm fixed effects specification which effectively compares the within-firm growth rate in the meetings groups to that in the control group. We estimate that by midline the log sales of treatment firms increased by a significant 0.078 more than that of control firms, corresponding to treatment effect on sales of 8.1 percent. The meetings also significantly increased profits, employment, and fixed assets. The magnitude and significance of these effects persisted to the endline survey. For example, baseline-to-endline change in log sales was 0.098 higher in treatment than in control firms ($p < 0.05$), corresponding to a long-term treatment effect on sales of 10.2 percent. These results show that the meetings had large and persistent positive effects on firm performance.

Turning to intermediate outcomes, we find that the meetings significantly and persistently increased the number of clients and the number of suppliers, as well as formal and informal borrowing. However, they did not significantly increase firms' tax-to-sales ratio. Using only the midline and endline survey—as data on it was not collected in the baseline—we also find that the meetings treatment had significantly and persistently improved management practices. Taken together, these results suggest two broad mechanisms at play: improved firm-to-firm matching which may have created new partnerships; and learning from peers which may have improved management. However the results do not conclusively identify these mechanisms: it is also possible that the meetings created growth through a different channel, and growth improved increased the demand for suppliers or improved management.

Both of these mechanisms suggest that group composition should matter: having better peers should generate better firm-to-firm matches and higher learning. Finding positive peer effects would thus add internal consistency to our results. To estimate peer effects—similarly to the approaches of Sacerdote (2001) and Shue (2013)—we ask whether firms randomized into groups in which peer firms

are larger grow faster. We find that having peer firms that employ more people leads to significantly higher sales, profits, productivity and number of clients. Using only data from the midline and endline survey, we also find that having larger peers leads to improved management practices. However, we do not find an effect on employment or the number of suppliers. Overall, these results are consistent with the view that meeting the managers of larger firms is more beneficial.

We next turn to additional interventions that generate direct evidence on channels. We begin with information diffusion, and establish three results about application rates for the financial products. (i) Informed managers in the meetings groups were significantly more likely to apply than informed managers in the control group. This result indicates a complementarity between information and the meetings; a possible explanation is that peers encourage the informed manager to apply. (ii) Uninformed managers in groups with a higher informed share were much more likely to apply. This is direct evidence for information diffusion. (iii) For the firm funding product—which is more rival because it can help a competitor’s business—diffusion was weaker in groups with higher competition. But for the (less rival) private savings product, the diffusion rate was not associated with the degree of competition. These results show that the meetings did diffuse business-relevant information, and that diffusion was weaker in the presence of competition.

Finally we identify the role of repeated interactions in shaping new partnerships and trust. We show that firms establish a significant 1.15 more direct partnerships—supplier, client, or joint venture—with their regular group members than with their cross-group members. Firms also get referrals from a significant 2.13 more peers in their regular group than in their cross-group. These results support the view that the meetings contributed to growth by reducing the cost of partnering; otherwise there would be no reason for the partners to come through the regular- and not the cross-group. We also find that firms exhibit significantly more trust in hypothetical trust games with their regular than with cross-group partners. A natural interpretation is that repeated meetings build trust which then helps create new partnerships, suggesting that in our context lack of trust is a key barrier to business connections.

At the end of Section 3 we discuss concerns with identification and interpretation. One issue is that experimenter demand effects may drive the results. While most of our data come from self-

reports in the surveys, for sales we also have book value because we asked managers to physically show us their sales value in the book. The difference between book sales and self-reported sales is small, insignificant, and does not vary with the meetings treatment, suggesting that demand effects in our main regression are small. Such effects are also unlikely to drive the result on peer effects which is identified using only managers in the meetings treatment. A second concern is that the meetings may have side-effects such as better access to government officials, and these may drive the results. But government officials were only involved in the first meeting, and they were also the ones who introduced us to the control firms. Moreover, this logic cannot explain the results on mechanisms that use only firms in the meetings. A third issue is collusion: perhaps firms in the meetings improved performance by coordinating price increases. But these firms are small actors in a large market. Also, standard models of collusion would predict that the increase in sales is accompanied by reduction in quantity, contradicting the positive employment effect of the treatment. And collusion cannot easily explain other gains, such as improved management or increased borrowing.

In the concluding Section 4 we discuss the magnitude, external validity and policy implications of the results. Business training interventions offer a natural benchmark to which to compare our estimates. Most training interventions have modest and often insignificant effects on firm performance (McKenzie and Woodruff 2014), though this likely partly because of power issues. The impacts of our intervention seem more robust, affect more outcomes and are more precisely estimated. Because our intervention is longer-term and fairly intensive, personalized business consulting is perhaps a better benchmark. Bloom, Eifert, Mahajan, McKenzie and Roberts (2013) find that intensive management consulting increased firm productivity by 17%. Our simpler treatment had a smaller but comparable productivity effect of about 8.9%. Concerning external validity, the evidence on the mechanisms suggests that the meetings help overcome information and trust frictions, which are likely to be important barriers for many young firms in developing countries. We thus expect that in such contexts organizing business associations can be an effective low-cost tool to foster private sector development.

Literature. A body of work highlights the importance of firm-to-firm interactions. Theories

include Acemoglu et al. (2012), Oberfield (2013) and Eaton, Kortum and Kramarz (2015), who build models of supply chains and study their aggregate and efficiency implications. Empirically, McMillan and Woodruff (1999) and Khwaja, Mian and Qamar (2011) show that interfirm relationships shape access to credit, and Shue (2013) documents that managerial networks influence compensation policies. In more recent work, Bernard et al. (2015) combine a search model of interfirm relationships with the natural experiment of the Shinkansen speed train in Japan to show that improvements in business travel increased the number of partners as well as firm performance. And in a pioneering experiment, Fafchamps and Quinn (2014) invite managers to serve as judges in business competitions, and use the resulting variation in networks to document limited diffusion of management practices. We contribute to this work with a design which is explicitly geared towards creating business links, identifies multiple mechanisms, and directly links to policy.

We also build on research that uses experiments to study private sector development. de Mel, McKenzie and Woodruff (2008) measure the return to capital in microenterprises, Bloom et al. (2013) and Bruhn, Karlan and Schoar (2013) measure the impact of management consultancy in different contexts, and McKenzie and Woodruff (2014) provide a review of some of this work. Our contribution is to evaluate a different intervention, that of business associations. And we also contribute to the literature that attempts to explain differences in productivity across firms (Syverson 2011) by showing that variation in business networks may be an important factor.

2 Context, experimental design and data

2.1 Context

Our experimental site is Nanchang, the capital city of Jiangxi province, which is located in south-eastern China. The city has a population of around 5 million people. In 2014, the GDP of Nanchang was 58 billion dollars, which ranked it as the 19th among the 32 capital cities in China. Nanchang is a fast-growing city 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.

2.2 Interventions

Basic experiment. In the summer of 2013, through CIIT we invited microenterprises and SMEs established in the preceding 3 years in Nanchang to participate in business associations. Around 5,400 firms expressed interest. We randomly selected 2,800 firms as our study sample. Out of this pool, we randomly selected 1,480 managers—the treatment group—and randomized them into meetings groups with 10 managers each. We organized the randomization as follows: first we divided the study area into local regions, and then we randomized firms into treatment and control, and randomized treatment firms into meetings groups, at the local region level. This design ensured that managers in the same meeting group did not have to travel far to meet each other. The 1,320 control firms were informed that there was no room for them in the meetings.

In August 2013, in collaboration with CIIT we organized the first business meeting for the meetings groups. In the first meeting only, we offered managers print material containing business-relevant information. We gave the same material to control firms as well. CIIT chose one of the managers in each meeting group as a group leader. This person was responsible for planning and scheduling all subsequent monthly meetings. In most groups, members took turns in hosting the meetings. We expected that managers in the meetings would discuss issues they face in running their business, learn from peers about business practices, establish new contacts, and more generally build social and business ties. For each meeting, the group leader took notes on the location, date, topics discussed, and the main takeaways, and submitted the log to us.

To provide incentives to participate, managers who answered our surveys and attended at least 10 out of the 12 monthly meetings got a certificate from CIIT. The certificate provides improved access to government services, such as government funding and admission to the MBA program offered by a local university. To firms randomized into the control group we also offered the certificate if they answered our surveys. We gave all firms the certificate after the conclusion of the one-year program.

Group composition. To measure mechanisms and explore the motives behind network formation,

we introduced additional interventions. First, to measure peer effects, we created variation in the composition of groups by size and industry. Almost all of our firms are in two broad industry categories, manufacturing and services, where services are primarily business services. In each region, we created two firm size categories, “small” and “large” by the median employment of firms in our sample in that region. We then created four kinds of groups in each region: small firms in the same industry; large firms in the same industry; mixed size firms in the same industry; mixed size and mixed industry. We randomized firms into these groups in each region.

Information treatment. To identify the role of networks in facilitating information transmission, we provided randomly chosen managers of both treatment and control firms with information about two financial products. The first is a government funding opportunity for the firm; the second is a savings opportunity for the manager. The funding product for the firm is attractive as it provides a non-refundable grant of up to RMB 200,0000 (about USD 32,000). This product is saliently in limited supply: each year only around 150 microenterprises or SMEs are selected to receive the funding. Because it can help a rival firm, managers may view this product as “competitive” and not share information on it with their competitors. The saving product is attractive because it offers an annual return of almost 7%, which is higher than the normal return of other saving products (about 5%). This product is also in limited supply because it people can invest in it only up to the point when the aggregate investment reaches RMB 5 million; but this rule is not very salient. This product may be viewed by subjects as “less competitive” because it is used by the manager, not by her or his business.

For each financial product, we distributed the information by phone calls and text messages to 0%, 50% or 80% of the managers in randomly selected business groups. Approximately one third of the meeting groups was assigned to each of the three levels of information treatment intensity.² We also distributed the information to 40% of control firms to ensure that the same share of treatment and control firms have the information. We independently randomized the information treatments for the two financial products.

Cross-group meetings. To learn about the roles of search costs and lack of trust as barriers to

²We stratified this randomization by group type.

building business connections, we organized one-time cross-group meetings. We took 439 managers in the meetings treatment and grouped them into 43 groups of around 10 managers such that no two managers in the same meetings group were in the same cross-group. These cross-groups met once in February of 2014. Moreover, in the 2014 midline survey we asked managers to play trust games (with large hypothetical payoffs) with a randomly selected regular group member as well as with a randomly selected cross-group member.

2.3 Surveys

We conducted a baseline survey before the intervention in 2013 summer, a midline survey after the intervention in 2014 summer, and an endline survey in 2015 summer. The surveys are conducted in person with the firm managers by our enumerators.

In the surveys we collect information from both treatment and control firms about the following groups of variables. (1) Firm characteristics. Profits, sales, costs, electricity use, spending on intermediate inputs, other balance-sheet measures, and innovation. (2) Managerial characteristics. Overconfidence, stress levels, happiness, and—in the midline and endline survey—questions on management style. The management questions covered five areas of management: evaluation and communication of employee performance, targets and responsibilities, attracting and incentivizing talent, process documentation and development, and delegation. (3) Firm networks. Business connections both within and outside the group, and the type of interaction (advice, referrals, purchases, sales). (4) Whether managers applied for the funding opportunities about which we had distributed information. (5) Employee satisfaction - only included in the endline survey.

2.4 Summary statistics and randomization checks

Table 1 shows summary statistics of firm and manager characteristics in the baseline sample. The table shows the mean for all firms, treatment firms, and control firms; and the final columns shows the difference between treatment and control firms. Panel A shows that the average age of firms in 2013 was between 2 and 3 years, and that almost all firms are private enterprises. Around 50% of the sample firms are manufacturing firms, with another 48% in the service sector. The average

Table 1: Summary Statistics: Firm and Manager Characteristics

	All Sample	Treatment	Control	Difference
<i>Number of Observation</i>	<i>2646</i>	<i>1409</i>	<i>1237</i>	
<i>Panel A: Firm Characteristics (2013 Baseline)</i>				
Firm Age	2.34 (1.75)	2.39 (1.72)	2.29 (1.77)	0.1 (0.068)
Ownership - Private non-SOE	0.98 (0.15)	0.98 (0.15)	0.98 (0.15)	0 (0.006)
Industry - Manufacturing	0.5 (0.01)	0.51 (0.013)	0.48 (0.014)	0.03 (0.019)
Industry - Service	0.48 (0.01)	0.49 (0.01)	0.47 (0.01)	0.02 (0.02)
Number of Employee	36.19 (86.49)	36.33 (90.63)	36.01 (81.55)	0.32 (3.37)
<i>Panel B: Managerial Characteristics (2013 Baseline)</i>				
Gender (1=Male, 0=Female)	0.84 (0.37)	0.846 (0.36)	0.837 (0.37)	0.01 (0.014)
Age	40.84 (8.85)	41.05 (8.46)	40.59 (9.27)	0.46 (0.34)
Education - College	0.29 (0.45)	0.288 (0.45)	0.295 (0.46)	-0.007 (0.018)
Government Working Experience	0.23 (0.42)	0.23 (0.42)	0.22 (0.41)	0.01 (0.02)
Communist Party Member (1=Yes, 0=No)	0.205 (0.4)	0.207 (0.4)	0.204 (0.4)	0.003 (0.016)
Working Hours - Weekday	9.62 (2.81)	9.64 (2.81)	9.6 (2.79)	0.04 (0.12)
Working Hours - Weekend	7.61 (4.53)	7.47 (4.57)	7.77 (4.48)	-0.3 (0.19)

Note :Standard deviations in parentheses for columns (1)-(3). Column (4) reports the difference in characteristics between treatment and control groups, and standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

number of employees is about 36.

Panel B presents statistics on managerial characteristics. 84% of managers are male, and in 2013 the average age of managers in our sample was around 41. 30% of managers have college education. We next look at indicators for government and political connections: 23% of managers worked in either government or state-owned enterprises in the past, and 20% of them were members of the party. Managers reported to work on average 9.6 hours during weekdays and 7.6 hours during weekends, suggesting that they were very busy and probably quite stressed. The fact that in spite of their intense schedules managers were willing to participate in the meetings suggests that they thought them to be valuable. There are no significant differences between treatment and control in any of the variables in the table, confirming that our randomization is valid.

Table 2: Summary Statistics: Business Activities

	All Sample	Treatment	Control	Difference
<i>Number of Observations</i>	2646	1409	1237	
<i>Panel A: Partnership (2013 Baseline)</i>				
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)	16.7 (20.3)	16.02 (17.94)	0.68 (0.75)
<i>Panel B: Borrowing (2013 Baseline)</i>				
Bank Loan (1=Yes, 0=No)	0.25 (0.43)	0.25 (0.44)	0.25 (0.43)	0 (0.017)
Informal Loan (1=Yes, 0=No)	0.12 (0.33)	0.114 (0.32)	0.13 (0.34)	-0.02 (0.013)
Other Loan (1=Yes, 0=No)	0.005 (0.073)	0.006 (0.075)	0.005 (0.07)	0.001 (0.003)
<i>Panel C: Accounting (2013 Baseline)</i>				
Sales (10,000 RMB)	1593.62 (6475.18)	1510.7 (5291.86)	1686.19 (7603.11)	-175.57 (252.32)
Log Sales (10,000 RMB)	5.59 (2.01)	5.6 (1.99)	5.58 (2.02)	0.02 (0.08)
Net Profit (10,000 RMB)	79.23 (205.35)	77.26 (199.92)	81.52 (211.55)	4.25 (8.09)
Log Valueadded/Number of Employees	1.97 (1.17)	2.00 (1.02)	1.94 (1.33)	0.06 (0.05)
<i>Panel D: Exit (2015 Endline)</i>				
Percentage of Firms Shut Down	10.92 (31.11)	10.86 (31.12)	10.99 (31.29)	-0.001 (0.012)

Note :Standard deviations in parentheses for columns (1)-(3). Column (4) reports the difference in characteristics between treatment and control groups, and standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.4 shows summary statistics on firms' business activities. Panels A and B present data on business connections with suppliers, clients, and lenders. The average firm reports to have had 46 clients and 16 suppliers. About 25% of firms borrowed from formal banks in the previous year, while 12% of firms have borrowed from friends and relatives. 1.3% of firms have loans from other sources such as private money lenders. The role of informal loans suggests that it is difficult for firms to borrow from banks, perhaps because it often requires collateral or guarantors from the government.

Panel C reports data on accounting measures. Average log sales was 5.6 and the average net profit was about 792,300 RMB (about USD 130,000). We use the valueadded to employment ratio as an indicator of productivity, and the average log productivity was 1.97. Finally, Panel D

shows that between the baseline and the endline survey about 10.92% of firms in our sample closed down. Consistent with the randomization, there are no significant differences between treatment and control in any of the variables in this table either.

The attrition rate from the baseline to endline survey was around 9.2%, not significantly different between the treatment and control sample. As a check on survey quality we also looked at the share of firms that did not change their answer between the baseline and the midline survey to some key questions where we expect variation over time. The share of firms for which sales, the number of clients and the number of suppliers did not change is 4%, 5% and 8%. These low shares suggest that basic errors caused by misreporting are unlikely.

3 Results

3.1 Effect of meetings on firm outcomes


Graphical evidence. We begin the analysis with graphical evidence that highlights some key patterns in the data. Figure 1 plots the kernel density of log sales for the treatment and the control group, both at baseline and at endline, i.e., both before the start and one year after the end of the 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 is—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 range of firm sizes.

Empirical strategy. To measure the impact of the meetings more precisely, and for inference, we now turn to estimate regressions. Our main specification is

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

Here i indexes firms, t indexes years, and y_{it} is an outcome variable such as log sales. The variable Meetings_{it} is an indicator for the treatment, which is time-invariant and equals one if the firm is

Figure 1: Kernel Density of log Sales



invited to the meetings. $Midline_{it}$ is an indicator for one year after the intervention, and $Endline_{it}$ is an indicator for two years after the intervention. 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 et al. (2008).

Our coefficients of interest are β_3 and β_4 , which measures—given the fixed effects specification—the differential change over time in the outcome variable in the meetings group versus in the control group. The key identification assumption is that firms in the meetings treatment do not have systematically different trajectories from those in the control treatment 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 will discuss such omitted variables in Section 3.5 below.

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. And, since our sample contains some larger firms, for specifications in which the dependent variable is neither binary nor a share between zero and one, we winsorize the regressions at 1% in both tails of the distribution.³

Results. We begin with Table 3 which focuses on a range of firm performance measures. Be-

³Not winsorized specifications yield similar results.

ginning with the effect one year after the treatment, column 1 shows that while in the control group firms' log sales increased insignificantly by 0.004, in the meetings treatment they increased by an additional significant 0.078, corresponding to a treatment effect on sales of 8.1%. Column 2 shows that average profits also increased by a significant RMB 257,500 (about \$36,000) more in the treatment group than in the control group. Columns 3 and 4 show evidence on factors of production. The treatment effect on log employment is a significant 0.052 corresponding to a 5.4% improvement in employment growth. The positive treatment effect on log total assets is also significant. Column 5 shows that the treatment had a positive but insignificant effect on log productivity of 0.05 ($p = 0.11$), corresponding to a productivity improvement of 7%. Finally, column 6 reports the treatment effect on the difference between (log) self-reported sales and 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. All of the above effects persists one year after the meetings program was over. For example, two years after we start the meetings intervention, the sales growth for treated firms was 9.8 log points higher than that of control firms. This suggests that the meetings intervention has a lasting positive impact on firm performance.

Table 3. Effect of Meetings on Firm Performance

Dependent var.:	log Sales	Profit (10,000 RMB)	log Number of Employees	log Total Assets	log Productivity	log Reported - log Book Sales
	(1)	(2)	(3)	(4)	(5)	(6)
Midline	0.00431 (0.0197)	11.89** (5.402)	0.0176 (0.0166)	0.0171 (0.0191)	0.0154 (0.0221)	0.00027 (0.007)
Endline	0.0134 (0.0286)	12.21 (8.278)	0.0286 (0.0243)	0.0404 (0.0345)	0.0390 (0.0349)	-0.0067 (0.0057)
Meetings*Midline	0.0780** (0.0359)	25.75** (12.59)	0.0524** (0.0264)	0.0645** (0.0308)	0.0502 (0.0386)	0.003 (0.0117)
Meetings*Endline	0.0984** (0.0497)	32.60* (18.52)	0.0768* (0.0438)	0.1156** (0.047)	0.0890 (0.0549)	0.00056 (0.0094)
Observations	7,857	7,664	7,857	7,851	7,537	7777
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.004	0.005	0.004	0.005	0.003	0.0016

Note: Standard errors clustered to the meeting group level for treated firms and to the firm level for control firms. Productivity is measured by the ratio between valueadded and number of employee. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We next turn to Table 4 which focuses on various intermediate outcomes that may have con-

Table 4. Effect of Meeting Treatment on Intermediate Outcomes

Dependent var.:	log Number of Clients	log Number of Suppliers	Bank Loan	Informal Loan	Tax/Sales	Stress
	(1)	(2)	(3)	(4)	(5)	(6)
Midline	0.0150 (0.0201)	0.0265 (0.0210)	-0.0396*** (0.0108)	0.0905*** (0.0113)	0.000592 (0.000976)	0.00820 (0.0194)
Endline	0.0442 (0.0290)	0.0487* (0.0286)	0.00769 (0.0136)	0.0507*** (0.0123)	0.00187 (0.00115)	0.0315 (0.0221)
Meetings*Midline	0.0898*** (0.0299)	0.0851*** (0.0308)	0.0907*** (0.0156)	0.0521*** (0.0175)	0.000728 (0.00149)	0.0444 (0.0277)
Meetings*Endline	0.118** (0.0459)	0.0897** (0.0406)	0.0790*** (0.0191)	0.0529*** (0.0170)	-0.00193 (0.00155)	0.0663** (0.0322)
Observations	7,841	7,826	7,857	7,857	7,849	7,857
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.011	0.009	0.014	0.042	0.001	0.005

Note: Standard errors clustered to the meeting group level for treated firms and to the firm level for control firms.

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

tributed to firm growth. Columns 1 and 2 show highly significant and persistent treatment effects on the log number of clients and suppliers ranging between 0.085 and 0.12. Columns 3 and 4 show that firms in the meetings treatment were significantly more likely to take out both formal and informal loans following the intervention. These results can be interpreted in two ways. One possibility is that the meetings help firms connect with more business partners and raise more capital, which then contributes to firm growth. An alternative is that the meetings generate growth through other mechanisms, which then translates into higher demand for business partners and for capital. In Section 3.4 we show direct evidence that the meetings did reduce the cost of firm-to-firm partnering. Column 5 shows that the tax-to-sales ratio of both treatment and control firms was essentially unchanged after the intervention. Thus the channel of the treatment effect is unlikely to be improvements in tax avoidance. Lastly, column 6 suggests that the meetings treatment increased managers' level of stress, perhaps because regular interactions, or business growth, created more pressure.

Finally 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 first standardizing and then averaging them. Because only the follow-up surveys contain data on

Table 5. Effect of Meetings on Firm Management

Dependent var.:	Management Score (Standardized)					
	Overall	Evaluation	Target	Incentive	Operation	Delegation
	(1)	(2)	(3)	(4)	(5)	(6)
Meetings*Midline	0.211*** (0.0509)	0.0938** (0.0458)	0.0338 (0.0428)	0.237*** (0.0472)	0.159*** (0.0502)	0.0710* (0.0409)
Meetings*Endline	0.215*** (0.0482)	0.0955** (0.0448)	0.0212 (0.0458)	0.223*** (0.0472)	0.179*** (0.0437)	0.0696 (0.0426)
Observations	5,211	5,211	5,211	5,211	5,211	5,211
Mid/Endline*Firm						
Demographics	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.011	0.002	0.000	0.013	0.007	0.001

Note: Standard errors clustered to the meeting group level for treated firms and to the firm level for control firms. Column (1) reports the impact of the meetings treatment on the overall management z-score, and columns (2)-(6) shows the intervention effect on five aspects of management: evaluation and communication of employee performance, targets and responsibilities, attracting and incentivizing talent, process documentation and development, and delegation. In all regressions, the interaction between midline/endline and firm demographics including firm size, industry, region, and their interactions have been controlled. *** p<0.01, ** p<0.05, * p<0.1.

management, we estimate the following specification:

$$y_i = const + \beta_1 \cdot Endline_{it} + \beta_2 \cdot Meetings_{it} \times Midline_{it} + \beta_3 \cdot Meetings_{it} \times Endline_{it} + Firm\ controls + \varepsilon_i. \quad (2)$$

Note that because the treatment itself is randomized, the coefficients of interest β_2 and β_3 are causally identified. The firm controls include the firm's region, size category (above or below the median employment in the region), industry (manufacturing or services), and all their interactions. Table 5 reports the results. In column 1, we estimate highly significant treatment effects of 0.21 at both midline and endline, measured in units of the cross-sectional 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 only exception being transparency to employees of firms' targets. Overall, we conclude that the meetings treatment had a large, persistent and highly significant positive effect on management practices.

Taken together, the results in Tables 3, 4 and 5 show that the meetings treatment substantially improved firm performance on several margins. We now turn to explore some potential underlying mechanisms.

3.2 Group composition and peer effects

As the first mechanism, we use the exogenous variation in group composition to explore how the identity of peers affects performance. Motivated by models such as Melitz (2003) in which productivity determines firm size, in our basic specification we measure peer quality with peer size. Using the sample of firms in the meetings groups, our starting point is the following specification:

$$y_{it} = \text{const} + \delta_1 \cdot \text{Midline}_{it} + \delta_2 \cdot \text{Endline}_{it} + \delta_3 \cdot \text{Midline}_{it} \times \log \text{Peer size}_{it} + \beta_4 \cdot \text{Endline}_{it} \times \log \text{Peer size}_{it} + \text{Firm f. e.} + \varepsilon_{it}. \quad (3)$$

Here Peer size_{it} is the average employment of the other firms in the meeting group of firm i at baseline, i.e., in the year before the intervention. The controls include the interactions of Midline_{it} and Endline_{it} with a set of firm demographics: indicators for region, industry categories at baseline (manufacturing or services), size categories at baseline (above or below the regional median employment), and all their interactions. As conditional on the firm demographics the groups—and in particular peer size—were randomized, in equation (4) the coefficients of $\text{Midline}_{it} \times \log \text{Peer size}_{it}$ and $\text{endline}_{it} \times \log \text{Peer size}_{it}$ are identified conditional on Midline_{it} or Endline_{it} interacted with these firm demographics.

Table 6. Effect of Peer Firm Size on Performance

Dependent var.:	log Sales	Profit (10,000 RMB)	log Number of Employees	log Productivity	log Number of Clients	log Number of Suppliers	Management
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Midline*log Peer Size	0.118*** (0.0322)	36.66*** (11.47)	0.0372 (0.0231)	0.117*** (0.0355)	0.0436* (0.0241)	-0.0300 (0.0302)	0.210*** (0.0352)
Endline*log Peer Size	0.166*** (0.0464)	52.25*** (17.27)	0.0149 (0.0375)	0.166*** (0.0429)	0.108*** (0.0343)	-0.0352 (0.0369)	0.111*** (0.0350)
Observations	4,183	4,076	4,183	4,029	4,173	4,170	2,774
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	No
Mid/Endline*Firm Demographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.040	0.040	0.070	0.055	0.068	0.051	0.159

Note: This Table is based on the sample of treated firms. Columns (1)-(6) are based on the three round survey, and column (7) is based on the midline and endline survey since the management section was not included in the baseline survey. Size of peers is calculated by the mean average of employment of group members. Firm demographics include firm size, industry, region, and their interactions. Standard errors clustered to the meeting group level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 6 reports the results. Column 1 shows that being randomized into a group in which peers employ on average 10% higher increases log sales by a significant 1.18% one year after the meetings

program and 1.66% two years afterwards. Column 2 shows that profits also increase by a significant RMB 36,660 (about USD 5,000). However log employment does not increase significantly. Column 4 suggests that been in groups with larger firms also improves firm productivity. Columns 5 and 6 show the results for suppliers and clients. Increasing average peer employment by 10% significantly increases the log number of clients by 0.43% in the midline and 1.08% in the endline. However, the point estimate for the log number of suppliers is negative, though small and insignificant.

We next turn to peer effects on management practices. Because the management data is only available in the midline survey, here we estimate

$$y_i = \text{const} + \delta_1 \cdot \text{Midline}_{it} \times \log \text{Peer size}_i + \delta_2 \cdot \text{Endline}_{it} \times \log \text{Peer size}_i + \text{controls} + \varepsilon_i. \quad (4)$$

Column 7 in Table 6 reports the results. The coefficient estimate of the midline coefficient (0.21) shows that having 10 percent larger peers results in having a significant 0.02 standard deviations increase in management practices. The effect shrinks by half but remains positive and significant one year after the end of the meetings intervention.

Why exactly does peer size matter? One possibility is that bigger firms have better management practices, which then diffuse to peers in the meetings. To explore this effect, we should measure peer quality with management practices. Since we had not included management questions in the baseline survey, we first construct a predicted measure of managerial practices for all firms at baseline. We do this by regressing, in the midline and endline control sample, the managerial index on log sales, log employment, manager education, age, gender, industry, and region indicators. We then use coefficients of that regression to compute the predicted managerial index at for all firms at baseline. We then estimate the same type of peer effects regressions, using this predicted value to construct a measure of peer quality. Table 7 shows that being randomized into groups with firms that have higher (predicted) managerial skills increases sales, profit, productivity, number of clients, and managerial ability.

Taken together, these results suggest that group composition matters: randomly assigned better peers generate more sales and profit, higher productivity, more clients, and better management practices.

Table 7. Effect of Peer Firm Management on Performance

Dependent var.:	log Sales (1)	Profit (10,000 RMB) (2)	log Number of Employees (3)	log Productivity (4)	log Number of Clients (5)	log Number of Suppliers (6)	Management (7)
Midline*Peer Management	0.147*** (0.0480)	51.62*** (17.66)	0.0298 (0.0354)	0.163*** (0.0535)	0.0800** (0.0362)	0.000800 (0.0404)	0.387*** (0.0481)
Endline*Peer Management	0.206*** (0.0701)	61.47** (26.23)	0.0412 (0.0529)	0.180*** (0.0629)	0.164*** (0.0501)	-0.00138 (0.0473)	0.144*** (0.0482)
Observations	4,143	4,037	4,143	3,990	4,135	4,132	2,748
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	No
Mid/Endline*Firm Demographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.037	0.039	0.069	0.053	0.070	0.052	0.163

Note: This Table is based on the sample of treated firms. Columns (1)-(6) are based on the three round survey, and column (7) is based on the midline and endline survey since the management section was not included in the baseline survey. Peer management score is calculated by the mean average of predicted baseline management scores of group members. Firm demographics include firm size, industry, region, and their interactions. Standard errors clustered to the meeting group level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

3.3 Information diffusion

We next consider the mechanism of the diffusion of business-relevant information. Here we exploit the intervention that to some randomly chosen managers we distributed information about two financial products: (i) a funding opportunity for the firm which managers likely consider competitive; (ii) a savings opportunity for the manager which managers likely do not consider very competitive. We randomized this information treatment independently across the two financial products, and provided the treatment to the same share of treatment and control firms.

Empirical strategy. We use two main regressions. First, using the full sample of treatment and control firms in the year after the intervention, we estimate, separately for each financial product:

$$Applied_i = const + \gamma_1 \cdot Info_i + \gamma_2 \cdot (1 - Info_i) \times Meetings_i + \gamma_3 \cdot Info_i \times Meetings_i + \varepsilon_i. \quad (5)$$

Here the dependent variable is an indicator for whether the manager reports in the midline survey to have applied for the product. The coefficient γ_1 measures whether the information treatment “works” in increasing the likelihood of application. The coefficient γ_2 measures whether uninformed managers in the meetings treatment are more likely to apply than uninformed managers in the control. This is a potential measure of information diffusion because it compares managers in the treatment who, on average, have informed peers, with managers in the control who likely do not. But it is an imperfect measure because the meetings treatment may affect applications not only through diffusion but also through firm growth. And γ_3 measures whether the effect of information

on applications is higher in the meetings treatment, i.e., whether the meetings complement the effect of getting information, for example because informed managers receive encouragement from group members.

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:

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

Here *Groupmember informed_i* is a dummy variable equal to one if at least one group member of firm *i* received the information treatment. Given that the information treatment is randomized, γ_4 measures the causal effect of having a higher share of informed group members on the decision to apply. *Competition_i* is an indicator for whether the meeting group of firm *i* has many competing firms. We define this variable by first computing the average number of in-group competitors of firms in each group (as reported by the firms themselves); and then splitting groups by the median of this value. Thus γ_5 measures the impact of higher competition on average diffusion, and γ_6 the extent to which competition reduces the strength of information diffusion.

The controls may include indicators for region, industry categories at baseline (manufacturing or services), size categories at baseline (above or below the regional median employment), and their interactions. We always include these variables when we include *Competition_i* in the regression. The reason is that the exogenous variation in group composition was created at the level of a region, and we formed groups based on industry and size categories. With the controls we are effectively comparing between firms who are in the same region, industry and size category. Because it was this pool of firms which we randomized into groups with homogeneous or heterogeneous members, by including the controls we are identifying the effect of *Competition_i* using only the exogenous random variation.

Results. Table 8 shows the results for the information treatment about a funding opportunity for the firm. The first two columns show the results from regression (5). In the first column we

Table 8. Diffusion of Information about Funding Opportunity for the Firm

Dependent var.:	Applied for the Firm Funding Product				
	(1)	(2)	(3)	(4)	(5)
Sample:	<i>All Firms</i>		<i>Uninformed Firms in Meetings</i>		
Info	0.300*** (0.0208)	0.370*** (0.0227)			
No Info * Meetings		0.202*** (0.0247)			
Info * Meetings		0.0721** (0.0323)			
Having Informed Group Members			0.315*** (0.0340)		0.402*** (0.0470)
Competition				-0.155*** (0.0497)	-0.0715** (0.0344)
Having Informed Group Members *Competition					-0.173*** (0.0605)
Firm Demographics	No	No	Yes	Yes	Yes
Observations	2,646	2,646	846	846	846

Note: This table is based on the midline survey. Competition equals one for groups with the average number of competitors (reported by firms) higher than median and zero otherwise. Firm demographics are indicators for firm size (above median employment in region at baseline), industry, region, and their interactions. Standard errors clustered to the meeting group level in parentheses. *** p<0.01, ** P<0.05, * P<0.1.

only include the $Info_i$ variable. The estimate shows that being informed increases the likelihood of application by a highly significant 30 percentage points, confirming that the information treatment worked. Column 2 also includes the interactions with the meetings treatment. Among uninformed managers, the meetings treatment increases application rates by a highly significant 20%. This can come either from information diffusion or from increased demand for funding because of firm growth. More surprisingly, among *informed* managers the meetings treatment also increases the probability of application by a significant 7 percentage points. This finding indicates that in our context formal funding and business networks complement each other, perhaps through group members' encouragement, or through increased demand for funds due to higher growth.

The remaining columns of the table report results from estimating variants of regression (6). The significant coefficient of 0.315 in column 3 shows that having at least one group member informed about the funding opportunity increases the probability of application by 31.5 percentage

Table 9. Diffusion of Information about Saving Opportunity for the Manager

Dependent var.:	Applied for the Private Saving Product				
	(1)	(2)	(3)	(4)	(5)
Sample:	<i>All Firms</i>		<i>Uninformed Firms in Meetings</i>		
Info	0.398*** (0.0182)	0.542*** (0.0232)			
No Info * Meetings		0.276*** (0.0276)			
Info * Meetings		0.00697 (0.0217)			
Having Informed Group Members			0.328*** (0.0310)		0.311*** (0.0462)
Competition				-0.00781 (0.0416)	-0.0224 (0.0380)
Having Informed Group Members *Competition					0.0456 (0.0615)
Firm Demographics	No	No	Yes	Yes	Yes
Observations	2,646	2,646	835	835	835

Note: This table is based on the midline survey. Competition equals one for groups with the average number of competitors (reported by firms) higher than median and zero otherwise. Firm demographics are indicators for firm size (above median employment in region at baseline), industry, region, and their interactions. Standard errors clustered to the meeting group level in parentheses. *** p<0.01, ** P<0.05, * P<0.1.

points. Column 4 shows that competition reduces application rates on average. And in column 5 in which we include $Groupmember\ informed_i$, $Competition_i$ and their interaction, we find a significant and large negative interaction effect of -0.173 , suggesting that competition may also reduce the strength of diffusion. Overall, these results show that the meetings channel information between managers, and that the dynamics of information diffusion is affected by competition.

Next, in Table 9 we turn to the information treatment about a savings opportunity for the manager. The structure is identical to that of the previous table. Column 1 shows that the information treatment was more effective for this product, and column 2 shows that complementarity between the information and the meetings treatments is stronger here. Column 3 presents very strong evidence for information diffusion, while columns 4 and 5 suggest that competition does not significantly reduce application rates or the strength of diffusion for this product. Overall, the stronger diffusion and the smaller competition effects are consistent with this product being both more popular and less rival than the other one.

Taken together, the results on the information intervention show that meetings improve the diffusion of business-relevant information. We also find that the extent to which information is rival affects the strength of diffusion. However—as in the model of Stein (2008)—even with rival information we find diffusion, suggesting that the benefits of sharing knowledge outweigh the costs of helping competitors. Thus improved access to information can plausibly contribute to the better performance of firms in the meetings groups.

3.4 Repeat interactions and new partnerships

We next turn to explore the role of repeat interactions by looking at the cross-group intervention. By comparing the number of partnerships in the regular groups and in the cross-groups, we can address two related questions: (i) whether new partnerships act as a mechanism through which meetings improve performance; and (ii) whether lack of trust is an important barrier to partnering.

To see the logic for (i), recall from Table 4 that firms in the meetings groups establish more partnerships. This result has two possible explanations. Either meetings reduce the cost of partnership and thus help growth, or meetings help growth through other channels, and growth increases the demand for partnerships. The cross-group treatment helps distinguish between these explanations: in the latter case we expect that firms establish the same number of new partners both in the regular- and in the cross-groups. For the logic behind (ii), note that managers may not organize meetings for themselves for at least two reasons. Either because there are search frictions in locating new partners; or because establishing new partnerships requires trust. Finding that the same number of new partners are created both in the regular- and in the cross-groups would be consistent with search costs being the main barrier, while finding more partnerships in the regular groups would be consistent with trust playing a role.

Table 10 uses the sample of managers in the cross-groups and shows measures of business relationships. Panel A reports the number of referrers—managers who referred suppliers, clients, partners, employees in different positions—in the regular and in the cross-groups. In the regular group on average 2.13 more managers act as referrers than in the cross-group, and this difference is highly significant. Panel B reports the number of direct business partners: suppliers, clients, and

Table 10. Repeated Interactions and New Partnerships

<i>Panel A</i>	Number of Referrers		Difference
	In Regular Group	In Cross Group	
Mean	2.18	0.06	2.13***
Standard Deviation	(0.083)	(0.62)	(0.079)
<i>Panel B</i>	Number of Direct Partners		Difference
	In Regular Group	In Cross Group	
Mean	1.44	0.29	1.15***
Standard Deviation	(1.49)	(1.52)	(0.07)
<i>Panel C</i>	Choice in Trust game		Dif
	In Regular Group	In Cross Group	
Mean	3.52	0.94	2.58***
Standard Deviation	(0.13)	(0.12)	(0.12)

Note: Referrers means a group member who has referred suppliers, clients, partners, managers, or employees to a firm; direct partner means a group member who is doing business with a firm. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

firms engaging in other joint business activities such as joint projects. There are a significant 1.15 more managers who act as direct partners in the regular group than in the cross-group. And Panel C reports average giving in hypothetical trust games played with a randomly chosen member of the regular group and of the cross-group. Managers exhibit significantly more trusting behavior towards their peers in the regular group.

These results imply that the meetings do reduce the cost of establishing new partnerships, so that partnering is indeed one of the channels through which they improve firm performance. Because we also find higher trust between managers in the meeting groups, a natural interpretation is that regular meetings create trust which in turn reduces the cost of partnerships. This logic parallels the findings of Feigenberg et al. (2013) that regular meetings build trust between borrowers and improve loan performance in microfinance. We conclude that lack of trust is likely to be an important barrier to creating business partnerships in our context.

3.5 Interpreting the results

Our estimates imply that the meetings have a large effect on firm performance. Here we consider potential confounds that may be driving this result.

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 over-reported their performance in the midline survey. Two facts suggest that demand effects are unlikely to drive our results. (i) Table 3 showed that the difference between the self-reported and the book value of sales does not vary with the treatment. (ii) Demand effects are unlikely to drive our results about mechanisms, which are identified from variation within the meetings treatment. Specifically, experimenter demand effects are unlikely to explain why the meetings improve the diffusion of business-relevant information, or the result that firms randomized with larger peers perform better. These findings show that—although demand effects may also be present—the meetings did have direct economic impact. We also note that all the treatment effects persists one year after the meetings concluded, while the experimenter demand effects should be weaken over time.

Outliers. Another concern is that some results may be driven by a few large firms, and the impact on the average firm is small. But Figure 1 shows that firms across a range of sizes were impacted. Moreover, we directly address this concern by winsorizing our main regressions at 1%. As the winsorized and the non-winsorized regressions yield similar results, we conclude that obvious outliers are not driving our findings.

Side-effects of the meetings. It is possible that the meetings improve firm growth not because of interactions between managers, but because of some kind of “side-effect”. One such side-effect is that firms in the meetings may have better access to the government through CIIT. Because—except for the first meeting—managers met without interference from CIIT or us, there is no obvious forum for regular access to CIIT officials. And since CIIT staff members introduced us to both the treatment and the control firms, it is not clear that treatment firms would have better government access than control firms. Thus the circumstances of the design make this effect unlikely. Moreover, this effect cannot easily explain some of the results on mechanisms, such as why being randomized with larger firms improves performance. Indeed, larger peer firms might actually crowd out the manager from accessing government officials. Still, to get at this mechanism more directly, using the endline survey we are analyzing data on firms’ business activities with the local government.

Another side-effect may be that firms in the meetings treatment can use a certificate that they

participate in the meetings to signal their quality, which brings them business. Importantly, control firms also get the certificate from CIIT, and all firms get the certificate after the meetings treatment is concluded. In the survey we asked for managers' willingness to pay for the certificate, and did not find a significant difference in that variable in either round of survey. Moreover, this logic cannot explain the results on information diffusion and on peer effects.

Collusion. Another issue is collusion: perhaps firms in the meetings improved performance by coordinating price increases. But these firms are small actors in a large market. Also, standard models of collusion would predict that the increase in sales is accompanied by reduction in quantity, contradicting the positive employment effect of the treatment. And collusion cannot easily explain other gains, such as improved management or increased borrowing.

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.

4 Conclusion

In this paper we used a field experiment with experimental business associations to measure the effect of business networks on firm performance. We found significant, robust, large, and persistent effects of the meetings on sales, profits, employment, productivity, and business partnerships. We also found evidence on several mechanisms: peer effects, information diffusion, and improved partnering and trust. And we argued that experimenter demand effects and other omitted variables are unlikely to explain all our results.

It is useful to compare our results to the impacts found in other types of interventions. McKenzie and Woodruff (2014) review several studies which evaluate business training and business consulting interventions. Concerning business training, one of their conclusions is that—perhaps because of power issues—most studies do not find a significant impact on sales or profits (see Table 9 in their paper). Among the exceptions are Calderon, Cunha and de Giorgi (2012) who find a 20% impact on sales and a 24% impact on profits; and de Mel, McKenzie and Woodruff (2012) who find a 41% increase in sales and a 43% increase in profits for start-up businesses. However, standard errors

are wide: out of these four estimates, only the de Mel et al. (2012) profits result is significant at the 5% level, and even for that estimate, we cannot reject a profit effect of as low as 6%. Our sales and profit impacts fall within the standard error bands of these estimates, but are more precisely estimated and significantly different from zero. McKenzie and Woodruff (2014) also review studies that use higher-intensity individualized consulting interventions on larger firms. The largest such intervention was conducted by Bloom et al. (2013) who introduced five months of free intensive management consulting to 11 in their sample of 17 large Indian textile firms. They found a productivity increase of 17%. The 8.9% productivity effect of our simpler intervention is smaller but comparable.

We conclude with a brief discussion of external validity. We emphasize that our meeting intervention was an intensive treatment: managers visited each others' firms, and spent hours discussing business issues. The results on meeting frequency suggest that the intensity of the intervention was important. Moreover, the results on trust and peer effects suggest both that trust barriers are important and that the meetings helped overcome it; thus in combination they are consistent with the view that managers viewed their peers as trusted experts and were willing to follow their advice. Based on these results we expect that similar business meetings should be useful in contexts in which trust and information frictions are important barriers. Because organizing the meetings is cheap, business associations may be an effective policy tool to foster private sector development in such settings. We are involved in a potential scale-up of our intervention, organized by CIIT, which may help us gather further evidence on the benefits of this policy.

References

- Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 2012, 80, 1977–2016.
- Bernard, Andrew B., Andreas Moxnes, and Yukiko U. Saito**, “Production Networks, Geography and Firm Performance,” NBER Working Papers 21082, National Bureau of Economic Research, Inc 2015.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts**, “Does Management Matter? Evidence from India,” *Quarterly Journal of Economics*, 2013, 128, 1–51.
- Bruhn, Miriam, Dean Karlan, and Antoinette Schoar**, “The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico,” Working paper 2013.
- Cai, Jing, Alain de Janvry, and Elisabeth Sadoulet**, “Social Networks and the Decision to Insure,” *American Economic Journal: Applied Economics*, 2015, 7 (2), 81–108.
- Calderon, Gabriela, Jesse Cunha, and Giacomo de Giorgi**, “Business literacy and development: Evidence from a Randomized Trial in Rural Mexico,” *Working paper*, 2012.
- de Mel, Suresh, David McKenzie, and Christopher Woodruff**, “Returns to Capital in Microenterprises: Evidence from a Field Experiment,” *The Quarterly Journal of Economics*, November 2008, 123 (4), 1329–1372.
- , —, and —, “Business Training and Female Enterprise Start-up, Growth, and Dynamics: Experimental evidence from Sri Lanka,” *Working paper*, 2012.
- 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*, 2003, 118.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz**, “Firm-to-Firm Trade: Imports, Exports, and the Labor Market,” *Working paper*, 2015.
- Fafchamps, Marcel and Simon Quinn**, “Networks and Manufacturing Firms in Africa: Results from a Randomized Field Experiment,” *Working Paper*, 2014.
- Feigenberg, Benjamin, Erica Field, and Rohini Pande**, “The Economic Returns to Social Interaction: Experimental Evidence from Microfinance,” *Review of Economic Studies*, 2013, 80, 1459–1483.
- Khwaja, Asim Ijaz, Atif Mian, and Abid Qamar**, “Bank Credit And Business Networks,” HKS Faculty Research Working Paper RWP11-017 2011.

- McKenzie, David and Christopher Woodruff**, “What Are We Learning from Business Training and Entrepreneurship Evaluations around the Developing World?,” *The World Bank Research Observer*, 2014, *29*, 48–82.
- McMillan, John and Christopher Woodruff**, “Interfirm Relationships and Informal Credit in Vietnam,” *Quarterly Journal of Economics*, 1999, *114*, 1285–1320.
- Melitz, Marc J.**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, *71* (6), 1695–1725.
- Oberfield, Ezra**, “Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture,” *Working paper*, 2013.
- Sacerdote, Bruce**, “Peer Effects with Random Assignment: Results for Dartmouth Roommates,” *Quarterly Journal of Economics*, 2001, *116*, 681–704.
- Shue, Kelly**, “Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers,” *Review of Financial Studies*, 2013, *26*, 1401–1442.
- Stein, Jeremy**, “Conversations Among Competitors,” *American Economic Review*, December 2008, *98*, 2150–2162.
- Syverson, Chad**, “What Determines Productivity?,” *Journal of Economic Literature*, 2011, *49*, 326–365.