

Experiments and Entrepreneurship in Developing Countries

Simon Quinn and Chris Woodruff*

University of Oxford

Abstract:

We discuss the value of experiments in illuminating constraints to the growth of firms in developing countries. Experiments have provided insight into both the value and difficulty of alleviating capital constraints in small firms. They suggest that urban, low-skilled labor markets appear to work reasonably well for firms, though there is a suggestion that frictions in markets for skilled workers may have more effect on firms. While observational data suggest that managerial training is important, experiments have shown that the traditional methods of delivering this training to small enterprises, at least, are not effective. Finally, while most work has focused on alleviating supply constraints, recent experiments have shown that positive demand shocks can be sufficient to generate firm growth. Experiments have been particularly illuminating in uncovering patterns in individual decision-making, showing how agents respond to the specific changes in circumstances or incentives generated by the experiment. They are most valuable when they complement insight driven by theory.

* Department of Economics and Department of International Development, respectively. Corresponding author: christopher.woodruff@qeh.ox.ac.uk. We thank Gabriel Ulyssea for discussions, Nicolas Pazos Navarro for research assistance and an anonymous referee for comments. Errors remain our own. This paper is based in part on a review prepared for the International Growth Centre with funding from the Argidius Foundation.

Experiments and Entrepreneurship in Developing Countries

1. New insights and new opportunities

In the shadow of Walmart de Mexico, a throng of street vendors typically gathers to sell food and goods from carts and baskets. The businesses of the large store and the vendors differ in almost every way. They use different mixes of capital and labor, they set prices differently, and they have very different management practices. But they co-exist within the same markets – even though their productivity, almost any way we measure it, is also vastly different one from the other.¹

Hsieh and Klenow's seminal 2009 paper builds a structure that rationalizes these types of productivity differences. In their framework, firms are differentially efficient because they face different factor prices. The prices they face are different because firms face different regulations or different market failures. Walmart pays labor and other taxes while the vendors are unlikely to do so. But Walmart also has access to sources of capital that the vendors lack, realizes scale economies on input purchases, has a global reputation for selling reliable products, realizes different markups, and so forth. We expect, in the absence of the regulatory and market constraints, more productive firms to grow and less efficient firms either to become more efficient or close down. But which prices are relevant for the continued coexistence of high- and low-productivity firms? Which policies need be altered so that the more efficient firms grow larger? In a model with two factors of production (capital and labor) but three potential sources of distortions (costs of the two inputs plus output distortions), we will not be able to isolate the policy-relevant source of distortions from the aggregate data.

Using a similar framework, Hsieh and Olken (2014) find that large firms have higher average (and, by the structure of their model, marginal) products of capital and labor than small firms do in Indonesia, India and Mexico. On the surface, this is striking because the experimental evidence on marginal returns to capital in very small firms, which we discuss in more detail shortly, suggests these returns are between 5 and 20 percent per month. Could the returns in large firms really be higher still? Unfortunately, it is not possible to infer directly from the Hsieh-Klenow or Hsieh-Olken frameworks the returns to relaxing access to capital for given types of firms. Large firms, for example, may have higher average (implied marginal) returns because they exercise market power. The framework simply says that larger firms are smaller than they should be given their productivity. Any policy, price, or circumstance that constrains their growth may will show up as high average / marginal returns to the factors included in the production function. Hence – though very valuable for pointing out the issue of dispersion in productivity – the framework is limited in its ability to guide policy.

Experiments in firms are relevant, then, because they allow us to measure directly the marginal effects of different types of constraints – and, therefore, to understand how changes in specific policies might improve productivity. Of course, academic researchers are not the only ones to experiment with firms: experimentation is the very essence of entrepreneurship itself (Kerr et al. 2014). High-growth entrepreneurs experiment at entry with product offerings, and continuously with product and process innovation. But researchers have joined

¹ Lagakos (2016) discusses the role of post-purchase transportation and storage costs in supporting the coexistence of “modern” and “traditional” retail technologies.

the experimental endeavor in force in the past decade, with an explosion in the volume of experimental work related to firms and markets.

In this paper, we review this recent academic literature through the lens of understanding how field experiments help illuminate constraints to firm and productivity growth. We aim both to summarize some interesting recent experimental results, and to speculate on how and where future work may provide important new insights. There are many ingredients in the recipe of a successful firm, and many factors potentially constraining firm growth. To guide our thinking, we organize the review around a production function familiar to every undergraduate economist: $Y = A \cdot f(K, L)$. We begin by discussing firms' access to factors of production: namely, access to capital and access to labor. We then move from the prosaic to the poetic: to consider not merely the individual factors, but the difficult entrepreneurial and technological tasks of combining them effectively. We take a broad view of what constitutes a firm: any non-farm enterprise, including the self-employed as well as larger firms. Indeed, we will see that much of the literature to date has focused on microenterprises.

Our review essentially makes two claims. First, that growth-oriented SMEs should be an important focus for private-sector policy in developing countries, because they represent important opportunities for future employment. Second, that this is an important time for experimental research on entrepreneurship in such firms – both because of the myriad new empirical insights in recent years, and because of the potential for exciting new research agendas in this area. In short, entrepreneurship matters for growth in developing countries – and field experiments are here to help.

2. Capital

Research related to firms has given the greatest attention to constraints on the ability of entrepreneurs to make optimal investments in fixed and variable capital. The largest part of the literature has focused on access to the external finance that funds capital investments in well-functioning financial markets. Entrepreneurs in low- and middle-income countries routinely list external finance as among the most important constraints limiting their growth. The seminal paper by Rajan and Zingales (1998) shows that capital-intensive sectors are under-represented in countries where financial markets are less developed. They show that in countries with less-developed financial markets, sectors that rely more on external finance grow more slowly *and* have lower rates of entry.²

Perhaps not surprisingly, then, access to finance and returns to capital investment were among the first enterprise-related topics to which experimental methods were applied in developing countries. As a result, the experimental literature on constraints to capital is more established and extensive than that of most other constraints. Based mainly on samples of very small firms, the literature demonstrates both the power and the limitations of the experimental approach.

2.1 Building capital through grants?

Several randomized experiments aim to provide a direct measure of the returns to capital by providing grants – either in cash or in kind – to small enterprises in developing countries. The

² Rajan and Zingales' sample is limited by the availability of data on accounting standards, and hence excludes the lowest-income countries.

results suggest that, on average, the return to such grants is high. De Mel et al. (2008) distribute grants of USD100 or USD200 to a random subset of a sample of microenterprises in Sri Lanka. The grants isolate a shock to the supply of capital that is independent of other characteristics of the firm. De Mel et al. use the capital shock to estimate that enterprises have average returns to the grant of five to six percent per month, representing annual returns of more than 60 percent. The marginal returns are almost twice as large in male-owned businesses, and near zero in female-owned businesses. They are also many times higher than microcredit lending rates. De Mel et al. (2012) report on a longer-term follow-up of the enterprises, showing that the returns on the initial grants are sustained five to six years later. Replications in Ghana (Fafchamps et al. 2014) and Mexico (McKenzie and Woodruff 2008) also measure marginal returns far in excess of microlending interest rates.

De Mel et al. also show that the experimental returns are more than twice as high as the returns from fixed and random effects estimations using the panel from the control sample in the experiment. Though there is bias in the non-experimental results from many sources, these results suggest that attenuation bias from mis-measured capital stock dominates. This highlights another advantage of experiments with firms: the relatively precisely measured marginal impacts induced by the experiment help to overcome measurement issues. These measurement issues are likely to be particularly acute in data from small firms.

Nevertheless, capital grants do not produce a measure of marginal returns to capital without additional assumptions. Owners may use their own resources to make complementary investments, or quickly de-capitalize the investment by, for example, drawing down inventories purchased with the grant. The capital injection may cause a change in labor input, and indeed, De Mel et al. show that owners respond by increasing the hours they work in the enterprise, at least in the short-run.

The indirect effects of capital shocks are one complication with experimental measures of returns to capital. Another limitation is that the experiments are more challenging with large firms. There is a more limited quasi-experimental literature on the returns to capital among larger enterprises.³ One notable entry in this literature is Banerjee and Duflo (2014), which uses changes in a lending program targeted to small and medium enterprises in India. Between 1998 and 2000, the Indian government changes the revenue thresholds defining eligibility for the set-aside program. These policy changes allow the researchers to identify credit supply shocks, and hence returns to borrowing among enterprises affected by the policy changes. The firms affected by the policy changes had capital stock of roughly USD 150,000 to 750,000, so are 2-3 orders of magnitude larger than typical firms reached by the cash grant experiments and microlenders. Using the shocks to access, Banerjee and Duflo estimate returns to capital of nearly 90 percent per year, substantially higher than the interest rates charged on the loans.

A second measure of the value of credit in larger firms comes from an experiment reported in McKenzie (2017b) in Nigeria. McKenzie designed the experiment around a business plan competition (YouWiN!) conducted by the Nigerian government with the support of the UK

³ There is also an extensive literature using observational data to estimate firm production functions, whether using the proxy variable approach (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Akerberg, Caves and Frazer, 2015) or dynamic panel methods (Arellano and Bond, 1991; Blundell and Bond, 1998). See Gandhi, Navarro and Rivers (2016) for a recent application of this approach in developing countries.

Department for International Development and the World Bank.⁴ The 475 entrants rated highest nationally or in their region were declared winners and awarded a grant of USD 50,000. Another 1,841 entrants rated just below this highest group were named “semi-finalists”, and 729 of these were then randomly selected to receive the grant. The random allocation of grants within the group of semi-finalists allows McKenzie to assess the impact of a substantial relaxation of the credit constraint, conditional on being good, but not great, in the eyes of the judges.

The grants are shown to increase survival and employment growth rates. Among entrants that were ongoing businesses (start-ups), the grant leads a 20 (37) percentage point increase in survival, and a 21 (23) percentage point increase in the likelihood of having 10 or more employees, three years later. McKenzie notes that measured profits are extremely noisy, and thus it is difficult to estimate the marginal returns on the grants. But the regressions for both new and on-going businesses show significant effects on profits only in the second follow-up survey, and then for both groups, the return is a bit less than 1 percent per month. It’s likely that these returns lag interest rates these sorts of firms would pay. Given the uniqueness of the very large grants provided, it is difficult to know how to generalize the results apart from the estimated returns to capital they generate. For example, the employment outcomes might have been smaller if the capital had been given as loans rather than grants.

2.2 Building capital through loans?

These experimental results on capital grants suggest optimism about the prospects for microcredit: after all, if returns to capital in microenterprises are high, then perhaps loans to those same enterprises would also generate high returns and encourage rapid growth.

Unfortunately, field experiments on microlending show little or no effect of loans on enterprise profitability or sales. In an important recent study, Meager (2018b) runs a Bayesian aggregation exercise on seven such experiments; she estimates “a precise and generalizable zero effect from the 5th to 75th quantiles”; for the right tail, she finds an imprecise positive effect, which is heterogeneous across contexts (see also Meager (2018a), on Bayesian aggregation of average effects). This is consistent with the conclusion of Banerjee, Karlan and Zinman (2015) – who, in summarizing the literature, “note a consistent pattern of modestly positive, but not transformative, effects”.

There are several important points to make about the design of these experiments, and what those designs tell us about how we should interpret the results. First, these experiments play on one of two margins. Either they randomize access at the neighborhood or village levels (Angelucci et al. (2015), Attanasio et al. (2015), Banerjee, Duflo et al. (2015), Crepon et al. (2015) and Tarozzi et al. (2015)), or they randomize marginal clients within all neighborhoods or lender branches (Karlan and Zinman (2011); Augsburg et al. (2015)) using credit scores to identify the relevant sample. Second, in either of these cases, we are learning about the effects on marginal borrowers rather than average borrowers.⁵ The design

⁴ For other recent empirical analyses of business plan competitions in developing countries, see Klinger and Schündeln (2011), Fafchamps and Quinn (2016) and Fafchamps and Woodruff (2017).

⁵ As Banerjee, Karlan and Zinman (2015) put it, “...the studies here identify impacts on marginal but not inframarginal borrowers... This is a strength in the sense that marginal borrowers are the focus of much theory, practice, and policy. But it is a weakness in the

determines the question the experiment is able to answer. For example, Karlan and Zinman sort applicants according to their credit score. Applicants scored 60 or above are all offered loans; those scored below 35 are all denied loans; and those scored between 35 and 59 are randomly sorted into a group that is offered a loan and a group that is not. This implies that we learn about the effect of loans for borrowers with credit scores between 35 and 59. We learn nothing about the effect of loans for borrowers with credit scores above 60 – where, plausibly, the effect of credit is different, and perhaps larger. Similarly, although Banerjee, Duflo et al. randomize at the neighborhood level, around 18 percent of households in their *control* neighborhoods receive a new microloan during the study period. The fact that the percentage in treatment neighborhoods (around 27 percent) is larger allows them to estimate the effect of microcredit. But the effect is estimated on the nine percent of households that are marginal borrowers – those who would not borrow in the control neighborhoods, but who do in the treatment neighborhoods. The effects on the sample affected by the experiment may be different from the average effect of microcredit across all borrowers.

The impact of microlending on enterprises is underwhelming in light of the high returns to capital found in experiments providing the capital shocks through grants. What might explain these differences? Where loans are made with the group-lending model, Fischer (2013) suggests that group members might pressure others in the group to make lower-risk investments that are also lower-return investments. But many microlenders, including some of those in the studies summarized by Banerjee, Karlan and Zinman (2015) now make individual loans. In this context, Field et al. (2013) suggest that borrowers may still make investments that are sub-optimally risky. They conduct an experiment with an MFI in which some borrowers are randomly chosen to receive an additional two months to make their first loan payment. Field et al. show that this grace period leads to the entrepreneurs making higher-return investments. But they also show that these investments are higher-risk and lead to much higher default rates by the borrowers. A structural model suggests that the grace period contract would not be profitable for the lender, even though the social returns to that contract exceed the social returns to the standard contract.

2.3 Building capital through equity investments?

Underlying the lack of appeal of the grace period loan for lenders is the fact that lenders do not capture the upside when the investments are successful, but suffer the downside when the investments fail. The obvious solution to this is some sort of equity contract that allows some risk (and reward) sharing by the investor. But there are challenges in microequity contracts given the lack of rigorous accounting and auditing standards and the lack of standard exit strategies for investors in these small firms. De Mel et al. (2018b) report on a failed micro-equity experiment in Sri Lanka. Several other micro-equity experiments are currently ongoing.⁶ There are many challenges in making small-scale equity investments, and small-scale experimentation will likely be key to gaining the knowledge required to design contracts in this space.

sense that impacts on inframarginal borrowers are key to understanding the totality of microcredit's success or failure as a development tool.”

⁶ Relatedly, there are active angel investor groups in some lower-income countries. But the networks of investors are informal – and, so far as we are aware, there is no analysis of their outcomes. See Kerr et al. (2014) for an analysis of the investment / mentoring effect of two angel investor groups in the U.S.

3. Labor

Workers acquire skills both in formal education and on the job. In well-functioning labor markets, skills will be reflected in wages; it is not surprising, then, that we see positive correlations between earnings (on the one hand) and education and labor market experience (on the other). But there is evidence that the labor market returns to both education and experience are lower in low-income countries compared with high-income countries. For example, Lagakos et al. (2018) compile data from labor surveys in 18 countries and show that the earning-experience profile is twice as steep in high-income countries as it is in low-income countries. Assuming that workers are paid according to their marginal product, this suggests that workers acquire skills on the job less quickly in lower-income countries.

There are at least three important mechanisms to explain the differences in returns to experience in low- and high-income countries. First, the workers themselves: it may be that a year's experience leads to the acquisition of fewer skills in low-income countries. Perhaps firms in low-income countries do not have the same incentives to provide on-the-job training as firms in high-income countries. Second, the firms: it may be that firms in developing countries are generally less productive (for example, through weaker management practices), and the consequences of this may be more pronounced as one moves up the skill/experience gradient. Thus, for example, an experienced production supervisor in Bangladesh may add less value to a garment manufacturer than an identical supervisor in a US firm.

Third, labor markets in low-income countries may be less effective at matching workers to firms – so, for example, our production supervisor may not have her particular skills recognized and put to good use (or, indeed, remunerated). Potentially, all three mechanisms are important constraints for firm performance. While any of these mechanisms might be suited to the use of field experiments, most of the recent experimental work has focused on the effectiveness of labor markets in matching workers to jobs. We therefore focus the discussion in this section on that mechanism.

The experimental literature has focused on matching of either unskilled or semi-skilled workers to firms.⁷ These are the segments most applicable to micro- and small-scale enterprises. Most of the research has taken the perspective of workers rather than firms, though there are some exceptions. Summarizing the research on active labor market policies more generally, McKenzie (2017a) concludes: “urban labor markets appear to work reasonably well in many cases, with fewer market failures than is often thought”.

In this section, we draw a distinction between (i) unskilled labor, (ii) semi-skilled labor, and (iii) managerial positions. Of course, this distinction is necessarily stylized; nonetheless, it serves to emphasize the central importance of labor market segmentation in developing countries. As our discussion makes clear, there are fundamental differences between the labor market experiences of those with different levels of skills; when it comes to labor market policy, we should not expect that one size fits all parts of a segmented market.

3.1 Unskilled labor

⁷ Several experimental papers have also explored other aspects of labor market search frictions from the perspective of the worker. In particular, Beaman and Magruder (2012) tests the role of referrals through a social networks experiment. (See also the more recent related work of Beaman, Keleher, and Magruder (2018) and Witte (2018).)

We should expect match frictions to be least important among unskilled workers, and the existing experimental evidence suggests this to be the case. In an experiment in Sri Lanka, De Mel et al. (2018a) ask whether microenterprises are constrained by access to labor. Starting with a sample of urban enterprises employing 0-2 paid workers, they offer a random subset of the firms an incentive to hire an additional worker. The incentive pays a fixed amount of 4,000 LKR (\$35) monthly for six months if the enterprise hires a fulltime worker from outside the family. Enterprises respond to the incentives: compared with the control-group enterprises, about one in seven treated firms hires an additional worker. However, within six months of the removal of the wage subsidies, there is no difference between employment in treatment and control enterprises. Moreover, De Mel et al. estimate that the increase in profits of the treated firms are not more than the amount of the subsidy. The results give little evidence of significant hiring frictions in this context.

Hardy and McCasland (2017) conduct a similar experiment in Ghana, intervening in the matching of enterprises to apprentices, a common entryway to employment. Apprentices in Ghana typically make an upfront payment to firms for their positions, and receive wages based on firm output during the apprenticeship period. Hardy and McCasland model the upfront payment as a screening mechanism, which works because low-ability workers know they will not recoup the fee with their (low) output-based wages.⁸ In the experiment, enterprises were randomly matched with zero to eight apprentices. The 35 percent of the firms assigned no apprentices in the program serve as a control group against which to measure the effects of apprentices on enterprise outcomes. Hardy and McCasland show that the matching treatment indeed increased employment in enterprises assigned apprentices, and that the additional employment generated increases in sales and profits during the two-year apprenticeship period. Since apprentices might be paid less than market wages, it is not clear how we think of these returns relative to regular wage workers, though the results clearly indicate imperfections in the apprenticeship market.

Experimental work on unskilled labor in larger firms has emphasized particularly the important role of turnover. Specifically, Blattman and Dercon (2018) conduct an experiment in five factories in Ethiopia that randomly selects workers hired for entry-level positions from among a large pool of applicants. They find that one-third of those offered the job quit within a month, and three-quarters quit within a year. Whether these turnover rates are the result of poor matches or are part of more complex employment strategies of workers is difficult to say based on the evidence in these studies. Note that the fact that Blattman and Dercon's firms are willing to randomize the selection of workers by itself suggests that the firms believe they face low search frictions. Large numbers of applicants gather outside the factory gates looking for entry-level positions. But the high turnover rates may suggest that the firms don't understand how they should be selecting workers. For example, employers may focus on the technical skills required for the position, while perhaps non-cognitive skills are more relevant.

3.2 Semi-skilled labor

⁸ Hardy and McCasland's experiment substitutes the monetary screening mechanism with a time-based screening mechanism, requiring youth interested in apprenticeships to attend a series of meetings.

Literature on semi-skilled labor has rightly placed substantial emphasis on both training and certification. A recent contribution in this space is the work of Alfonsi et al. (2017), who run an experiment among unemployed youth in Uganda, comparing vocational training with wage subsidies designed to induce on-the-job training. Importantly for the present review, one arm of the project randomizes the number of workers assigned to firms, allowing an estimate of the returns to the firms of hiring these workers. Consistent with de Mel et al. Alfonsi et al. find little evidence of significant frictions in the matching of firms to workers. Absent the subsidies, only 13% (19%) of firms agree even to meet with the vocational trained workers (unskilled workers) matched to them by the researchers. They conclude that “there is not much evidence for search frictions related to meeting untrained workers or meeting skilled workers in these labor markets.”

The Ugandan firms receiving incentive to train apprentices retain the majority of these workers after the subsidy ends. But consistent with the lack of frictions in the matching market, essentially all of the workers leave within three years and there is full displacement – the treated firms are no larger than control firms two years after the experiment. However, the treated firms are significantly more profitable over the post-treatment period, and the wage subsidies can account for only one-third of the increase in profitability. This implies either that the subsidized firms hire more productive workers than these firms typically hire, or that the subsidized firms retained a higher share of workers’ output than they typically would. These changes are worth understanding in more detail, but the central result here is that there is no evidence that firms are larger as a result of the increased incentives to hire workers.⁹

Alfonsi et al. is unusual in directly measuring the effect of matching on firms’ growth. More typically, studies report effects on program participants without measuring general equilibrium effects outside the sample. Nevertheless, the results on other experiments provide at least suggestive evidence that matching frictions have effects on firm outcomes. For example, Abebe, Caria et al. (2018), report results on an experiment providing job applicants transportation subsidies and a job application workshop (which included a skill certification component) among educated youth in Addis Ababa. They find short-run effects on the quality of employment from both interventions. Four years later, those who received the certification treatment have wages approximately 20 percent higher than the control group. They are also significantly more likely to say that their current job makes regular use of abilities acquired in either previous jobs or at school. If wages track marginal products, these results provide suggestive evidence that certification likely produced gains for the hiring firms as well. Abebe, Caria et al. (2018) conduct a separate experiment that compensates job applicants for the transportation and other costs of applying for a job. They show that firms believed that the application costs act as a screening device (similar to the mechanism in Hardy and McCasland) and therefore, that the subsidies would lower the quality of the applicant pool. Abebe, Caria et al. show that the firms’ beliefs are incorrect, and that the subsidies *increase* the quality of the applicant pool. These results suggest that more direct evidence of the effect of matching frictions on firms hiring more skilled workers is a promising direction for experiments to go.

3.3 Managers

⁹ Alfonsi et al. find that the vocationally-trained workers enjoy greater skill transferability – and, therefore, job mobility. See also Bassi and Nansamba (2017).

Experimental work to date has focused primarily on matching frictions in unskilled or semi-skilled labor markets. Labor markets for managers may have much larger consequences for firms that grow to the point that hierarchies become relevant. This is, of course, a challenging area for experimental research – both because firms large enough to have managerial hierarchies are fewer in number and because the larger firms are generally less able to find time for researchers. We review the experimental work on training managers below, but we are unaware of any experimental work on selection of managers.

A novel (non-experimental) approach studies the way that different CEOs divide their time. Bandiera et al. (2017) collect time use data from 1,114 CEOs of firms in six countries (Brazil, France, Germany, India, the UK and the US). Using these data and a machine learning algorithm, they classify CEOs as “managers” or “leaders”. Managers involve themselves in the details of production while leaders have group meetings and delegate. Bandiera et al. show that leader CEOs run firms that are larger and more productive than the firms run by manager CEOs. Moreover, using a subsample of firms that change CEOs during the period covered by their productivity data, they show that performance of a firm improves after it hires a leader-type (rather than a manager-type) CEO. This leads them to ask whether firms make mistakes in hiring managers rather than leaders and if so, why they do so. Mapping their data to a simple model leads to the conclusion that firms are heterogeneous in needs – some firms benefit from manager-type CEOs, but matching frictions imply that even some firms that would benefit from leader-type CEOs end up with manager-types. This mismatching is much more prevalent in the lower-income countries (Brazil and India) in their sample.

One explanation for the greater number of mismatches in lower-income countries is that matching frictions are higher where trust is lower. A lack of trust may cause owners to limit the pool from which they select managers – in the extreme, for example, to family members (see Lemos and Scur, 2018). Moreover, leader-type CEOs must delegate detailed decision-making to subordinates. This delegation implies a level of trust, either in the individuals per se, or in the power of the incentives those individuals face. Bloom, Saddun and Van Reenen (2012) show that firms located in higher trust regions, or multinationals headquartered in higher-trust countries, have more decentralized decision-making processes and are therefore able to grow larger.¹⁰ The market for managers has important implications for firms in developing countries. The non-experimental literature points in potentially relevant directions for experiments, though experiments on matching (as opposed to training, to which we return below) entail significant challenges.

4. The technology of firm performance

In previous sections, we have discussed constraints to investing in capital (‘K’), and to hiring labor (‘L’). We now turn to discuss firm technology: ‘A’. As every undergraduate economist knows, the concept of ‘technology’ in a production function can be a frustratingly elusive idea: the notion can potentially include everything affecting firm output other than

¹⁰ LaPorta et al. (1997) show that the 25 largest publicly traded firms are larger in countries with higher levels of trust, measured by responses to the World Value Survey generalized trust question. Leaven and Woodruff (2007) show that Mexican firms with better contract enforcement are larger. Boehm and Oberfield (2018) show that firms in Indian states where courts function better are larger and less vertically integrated, and Acemoglu, Johnson and Mitton (2009) show a similar pattern using cross-country data.

capital and labor. In the context of firms in developing countries, the term often refers implicitly to the individual owner or manager of the firm: does this person have a good business idea, and is this person good at managing others? For this reason, empirical literature in this space tends to focus on both the discovery of high-potential entrepreneurs (that is, “where are they?”), and on the processes by which such entrepreneurs can improve their skills (“how can training help?”). Broadly, this distinction accords with traditional views of how innovation happens in a market economy – namely, (i) by a Schumpeterian process of ‘creative destruction’, through which the market selects promising ideas, and (ii) through a process internal to the firm, by which individual managers make important decisions about the firm’s adoption of new processes. These are the questions that we pursue in this section.

4.1 Searching for Jobs...Steve Jobs

Firm performance in developing countries is typically extremely heterogeneous – with large productivity differences between the leaders and the laggards (see, for example, McCaig and Pavcnik’s (2016) descriptive evidence from Vietnam). Some of this difference is surely due to circumstance – some firms have stumbled upon more productive business opportunities, or are fortunate to have ready access to inputs. But much of the difference is innate to the firm managers themselves: whether through formal skills or informal know-how, some managers are much better at their job than others. Both from formal estimation and from informal observation, it appears that the proportion of firms with strong growth potential is small. Using data from Brazil, for example, Ulyssea (2018) estimates that 91 percent of unregistered firms are either subsistence or “rationally informal” firms, with little prospect for dynamic growth: only nine percent are constrained by informality.¹¹

How, then, can we find the small proportion of firms with high growth potential? Predicting which businesses will succeed is challenging even for highly skilled venture capitalists who have strong incentives. However, the sustained success of certain venture funds suggests that some experts are able to do better than random guessing. As we noted earlier, business plan competitions are a common means of promoting entrepreneurship and encouraging entrepreneurs to turn ideas into firms. But are the competitions successful at selecting firms with greater growth potential and, if so, can we learn anything about selecting those individuals at a lower cost?

One approach to this question is to ask if standard survey information can predict higher-performing firms – and, if so, whether the views of competition judges adds to these predictions. For example, Kahneman and Klein (2009) examine studies comparing predictions by experts with predictions based on simple combinations of quantifiable variables. They conclude that the predictions of experts are likely to be most valuable when (i) outcomes are reasonably predictable, (ii) experts have extensive experience making similar judgments, and (iii) the experts receive immediate feedback on the accuracy of their predictions when they make them. On the surface, none of these conditions hold when trying to predict which enterprises are likely to grow fastest – or, indeed, benefit the most from an intervention. These outcomes are highly variable, and usually take quite a long time to be

¹¹ This proportion is consistent with the experimental work reported in De Mel et al. (2013), who provide informal owners monetary incentives to register their business formally. Using the treatment as an instrument for formal status, the researchers find that being formal has an effect only for around five percent of the informal firms.

realized. Hence, we should expect that predictions from baseline measures might be more accurate.

With this exercise in mind, both Fafchamps and Woodruff (2017) and McKenzie and Sansone (2017) compare the ability of business plan competition judges and baseline survey data to predict firm growth. Fafchamps and Woodruff use data collected through a business plan competition with 335 applicants running on-going businesses in Ghana. Follow-up survey data allows them to measure survival and growth outcomes 12 and 24 months after the competition. They collapse the baseline survey data into three measures capturing ability, previous borrowing, and management practices, showing that ability is most strongly correlated with future outcomes. These three measures outperform expert options in an econometric horse race, though Fafchamps and Woodruff show that the rating of the judges adds predictive power even after controlling for the survey measures. McKenzie and Sansone carry out a similar analysis using data from the YouWin! competition, with a larger sample and follow-up data from three years after the competition. They also find that baseline data outperform the expert panels. However, in their data the panel ratings are uncorrelated with growth outcomes. We note that judges had live interviews with the entrepreneurs in the Fafchamps and Woodruff competition, but did not in the YouWin! competition. In both of these projects, one is left with the view that neither the surveys nor the judges can predict much of the variation in future growth.

An alternative approach to experts is to ask whether peers have information about which businesses are most likely to succeed, and if so, whether we can extract that information in an unbiased manner. Tapping into peer networks, of course, has a long heritage in development entrepreneurship. The approach was integral to the group lending model developed by Muhammad Yunus and Grameen Bank (Yunus, 1989). Hussam et al. (2017) carry out a project in Amravati, India following in this tradition. They work with samples of subsistence businesses, but their approach may have applications to more dynamic firms as well. Hussam et al. begin by dividing their sample into groups of five owners. They then ask each owner about other members of their group: their education level, enterprise profits and, most importantly, the marginal returns to an investment of USD 100 in their business. The design incorporates random cash grants that allow the researchers to estimate the actual marginal returns to capital in the sample. Although it is not obvious we should expect peers to be able to predict marginal returns of an enterprise, Hussam et al. show clear evidence that they are able to do so: enterprises ranked in the lowest tercile of expected returns show no gain in profits after receiving the cash grant; those ranked in the top tercile show increases in monthly profits of more than 20 percent of the grant. They also show that extracting information when stakes are high is challenging: when peers know their evaluations will affect who receives the grants, they are biased in favor of family and friends. Hussam et al. implement incentives for truthful reporting that they show, under certain conditions, do produce more accurate reports. Whether the conditions necessary for truthful reporting can be met in actual practice is unclear, but the work at least shows that peers have valuable information.

As in Fafchamps and Woodruff (2017) and McKenzie and Sansone (2017), Hussam et al. compare the prediction of peers with predictions based on “hard” data from baseline surveys conducted with the entrepreneur herself. They compare predictions from machine learning models with the predictions of peers. They find that machine-learning predictions can also help to isolate the owners with higher marginal investment returns. Those in the highest tercile of predicted returns realize monthly returns of 18 percent in the grants experiment. But the soft information of peers has predictive power even after controlling for the hard survey

responses. Again, the peers have sustained personal interaction with those they are judging, suggesting that this may be important to the ability of judges to add value.

Psychometric measures are an alternative approach to predicting outcomes. Dlugosch et al. (2018) use data from banks making small-scale loans in Nairobi to show that personality measures – conscientiousness, extroversion and integrity – predict loan repayments. However, a common theme in the literature on selection on growth outcomes is that entrepreneurial ability is the better predictor. Ability is typically measured by some combination of fluid reasoning tests (e.g., Raven non-verbal reasoning tests), numeracy, years of schooling, and similar measures. De Mel et al. find that higher-ability owners have higher marginal returns, and both Fafchamps and Woodruff (2017) and McKenzie and Sansone (2017) find that measures of ability predict growth. Fafchamps and Woodruff (2017) find no effect of attitude measures while Hussam et al. (2017) find that several psychometric measures only weakly predict returns to capital in their experiment. The fact that panels and peers (who are likely to incorporate soft information on attitudes) can predict outcomes when they have personal interaction with entrepreneurs (as in Fafchamps and Woodruff and Hussam et al) may suggest that we simply have not properly quantified the attitudes that matter for growth.

4.2 Entrepreneurship training and mentoring

If it is challenging to select entrepreneurs who are more likely to succeed, can we instead create entrepreneurs who are likely to succeed? Attempts to build entrepreneurial skills have taken different routes. In this section, we discuss entrepreneurship training programs; we then go on to discuss research on the effectiveness of individualized consulting and mentoring.

McKenzie and Woodruff (2013) review the evidence on the effect of training on performance of SMEs in developing countries. Reviewing 16 studies (mostly experiments), they find little evidence that the standard training models lead to improved performance of enterprises. A majority of the 16 studies use standard classroom-based training programs like the ILO's SIYB program on topics of record keeping and marketing. They show that most of the individual studies are statistically underpowered, particularly on key outcomes like profits that are integral to cost-benefit calculations. The lack of statistical power comes from the combination of highly heterogeneous samples used in typical studies and high levels of variation in the key outcomes, rather than small sample sizes *per se*. However, given the substantial number of studies reviewed, the results have generally been interpreted as indicating that standard training programs have at most modest effects on enterprise performance.

The underwhelming performance of standard training programs does not imply that training itself does not matter. Indeed, using data from seven countries, McKenzie and Woodruff (2017) show that the business practices that are the focus of the standard training programs are themselves highly correlated with differences in firm performance both in the cross section and in panel data. Firms implementing a larger share of 30 business practices relevant to microenterprises are larger, more profitable, more likely to grow over time, and less likely to exit. These basic results are based on observational data – though the correlation between practices and enterprise performance is robust to including a large number of controls for owner ability (which, not surprisingly, is itself related to performance). But a key factor is that even if the management practices affect performance, the enterprise owners

participating in training sessions implement very few additional best practices. Indeed, McKenzie and Woodruff show that training programs that have larger effects on management practices (for example, Anderson-MacDonald et al. 2017) also have larger effects on enterprise performance. In other words, the effectiveness of training delivery matters, and research suggests that most entrepreneurship training programs fail because they fail to improve business practices.

Why might firms choose not to implement the practices they are taught? One possibility is that, just as “every unhappy family is unhappy in its own way” (as Tolstoy famously put it in *Anna Karenina*), perhaps every firm faces its own set of challenges, each requiring its own individualized solution. That is, the process of maximizing profit – something that is often treated as a trivial technical problem in many models of the firm – may be deeply challenging, and may depend substantially on idiosyncratic circumstances of each firm. This idea – that economic knowledge is “concrete and contextual and not abstract” (Boettke, 2002) – has a long intellectual tradition in Austrian economics. In his 1945 essay on “The Use of Knowledge in Society”, Hayek distinguished “knowledge of general rules” with “knowledge of the particular circumstances of time and place”. The relevance of this insight for entrepreneurship training is spelled out in Hayek’s subsequent work – where Hayek went further, emphasizing that specific knowledge may be difficult, or even impossible, to articulate to others (Oğuz, 2010). In *The Fatal Conceit*, for example, Hayek linked his theory of knowledge to the problem of teaching entrepreneurial skills – writing that “so much knowledge of particular circumstances is unarticulated, and hardly even articulable (*for example, an entrepreneur’s hunch that a new product might be successful*) that it would prove impossible to make it public” (1988, page 89, emphasis added).

Similar themes have more recent resonance in organizational economics – in particular, with an emphasis on how heterogeneous firms are likely to face heterogeneous constraints to the adoption of potentially beneficial management practices. For example, Gibbons and Henderson (2012) describe “four ‘tions” (originally due to Jan Rivkin; emphasis in original):

First, managers may have problems of *perception* – they do not know they are behind. Second, managers may have problems of *inspiration* – they know they’re behind, but they don’t know what to do about it. Third, managers may have problems of *motivation* – they know they’re behind and they know what to do, but they lack incentive to adopt new practices. Fourth, managers may have problems of *implementation* – they know they’re behind, they know what to do, and they’re trying hard to do it, but they nonetheless cannot get the organization to get it done.

If we accept this premise – that different firms might have very different reasons for refusing to adopt improved management practices – then it should hardly surprise us that management training focused on generic skills is not particularly useful for many firms.

If generalized entrepreneurship training seems to lack promise, a natural question is whether individualized consulting programs can be more effective. A recent “instant classic” on this question is Bloom et al.’s (2013) experiment with large textile factories in Mumbai, India. A control sample of six firms received a one-month consulting intervention focused on record-keeping, to retain their interest and to ensure the quality and comparability of the information reported. A treatment sample of 11 received an additional four months of consulting services aimed at improving management practices. Using high-frequency data on output and quality defects, Bloom et al. show that the intensive consulting intervention significantly increased output per worker, quality, and TFP, and significantly reduced investment in inventory. Management practices, measured by 39 practices specific to textile production, improved.

Why did managers not adopt the practices on their own? There are both proximate and deeper answers to this question. Interviews with managers reveal that the proximate reasons are either that the managers were unaware of the practice (most often for “uncommon” practices), or that they did not think that adopting the practices would be profitable (most often for “common” practices). Of course, one can then ask why the managers did not bother to become better informed. We speculate on these deeper causes below.

A second major study in this area is Bruhn et al. (2018), who conduct an RCT with 432 firms in Puebla, Mexico to evaluate the effect of consulting and mentoring on firm growth. The 150 firms in their treatment group are much smaller than those in the Bloom et al. study. Around 70 percent of the treated sample (108 of the 150) are classified as “micro” (and hence, presumably have fewer than five employees), while 34 are “small” and eight are “medium”. The treated firms received weekly four-hour sessions over a period of one year. The researchers find generally positive impacts on profitability and return on assets at the end of the year of consulting, though these results are somewhat fragile (in that they depend on the specification). The more impressive results come from data taken from the national Social Security system for as long as five years after treatment. These data show that, compared with firms in the control group, those receiving consulting services grew faster each year after the program, leading to a 57 percent growth in employment after five years,¹² representing 5.7 extra employees in each treated firm. The consulting appears to have led most directly to improvements in marketing and recordkeeping.

Although effective in terms of growth, the cost-effectiveness in Bruhn et al.’s study is somewhat difficult to assess given the data available. The researchers estimate the cost of consulting at almost \$12,000 per firm, with 70-90 percent of that cost paid by the government. At 5.7 jobs per treated firm, that suggests a cost per job of around \$2,000, which compares quite favorably with other programs aimed at creating jobs.¹³ While they lack long-term data on profits, they note that the profit per worker would not need to be large for the intervention to pay for itself. The difficulty in interpretation comes from the fact that the follow-up data are averages for the treatment and controls groups. Baseline firm size is highly right-skewed, with a sample of 108 microenterprises and eight enterprises at least an order of magnitude larger. It seems likely the majority of the microenterprises in the sample would not find a \$12,000 consulting intervention viable were they to have to pay the full cost – but the data don’t allow us to judge that supposition one way or the other.

While there is evidence that individualized consulting can be effective, then, the model is expensive. One response to this is to attempt to simplify and lower the cost of management interventions, and there are several ongoing projects focused on models to deliver similar

¹² Only 57% of the enterprises in the sample were registered with the Social Security Institute (IMSS). For confidentiality reasons, the researchers did not receive individual firm-level data from the IMSS records, but rather received aggregate treatment and control group means and standard deviations. This is of some concern because, given the variance in the size of firms in the sample, the employment results could be driven entirely by a large expansion of a single medium-sized firm.

¹³ For comparison, McKenzie (2017) shows that the average cost of creating a job through nine vocational training programs assessed with experiments is \$17,000 - \$60,000. McKenzie (2016) estimates the cost of creating a job with the grants given to winners of the YouWin! Competition in Nigeria is around \$10,000. In their ‘Aspire’ business plan competition, Fafchamps and Quinn (2016) estimate a cost of \$1,250 per job created.

outcomes at lower cost. Note that Bruhn et al. show – as our earlier discussion suggested – that the changes the firms make after receiving the consulting services are both idiosyncratic and complex, so it is not obvious that these attempts at simplification will work.

A key question in this literature is how to fund the training. Although Bloom et al. show that the returns to very intensive management consulting are reasonably high, the market for these services has been slow to develop: firms appear reluctant to commit to paying for the training even though the large benefits that researchers have quantified are internalized by the recipient firm. Why? Rivkin’s “four ‘tions” can be useful here again. It may be that many firm managers simply don’t realize that they have a problem (that is, ‘*perception*’), or managers may doubt that consulting services – even individualized consulting services – can really be useful for their particular context (a problem of *inspiration*). Many academics would be optimistic that the lessons from Mumbai have implications in other countries and industries. But is it unreasonable for the owner of a scrap metal factory in Kenya (for example) to ask whether the evidence from Mumbai is relevant for him? Managers may also lack the *motivation* to change: we typically think of managers as seeking to maximize firm profits, but it is entirely possible that many managers resent the loss of control that outside consulting services imply. Bloom and Van Reenen (2007) find that firms facing less competitive markets and family-owned firms have poorer management practices. Lemos and Scur (2018) follow this by presenting quasi-experimental evidence showing that dynastic successions in family firms reduce the probability of adopting managerial best practices. Finally, firms may face problems of *implementation*. Though not management practices *per se*, Atkin et al. (2017) find that workers paid piece rates may misinform management about the value of technology when they believe that the transition to the technology involves costs for which they will not be compensated. Macchiavello et al. (2018) also find evidence that resistance from workers increases costs of adopting new practices, this time male workers in Bangladeshi garment factories experimenting with promotion of females to supervisory roles.

For these reasons, it is fair to ask whether and how knowledge generation will ultimately lead to firms engaging consultants more. More generally, there are many open empirical questions about the nature of entrepreneurial knowledge, and the means by which it can be effectively transmitted in different settings. To this end, recent literature has started to test novel variations on the training/consulting theme. We highlight here three novel approaches. First, Campos et al. (2017) report on experiments with microenterprise owners in Togo comparing the effect of standard training with ‘personal initiative’ training; the authors find that the ‘personal initiative’ training increases profits by 30% over two years, whereas the standard training has no significant effect. Second, Anderson-MacDonald et al. (2017) link small business owners in Uganda with mentors around the world via meetings over Skype. The mentors are typically MBAs and consultants. This is interesting because the mentors are likely to be very informed about good business practices generally, but may have little understanding of the Ugandan context. Hence, the project helps to inform us about the importance of this local context. Anderson-MacDonald et al. find that the Skype mentoring does not improve business practices, but may induce the Ugandan firms to “pivot” their business, for example, by stopping the production of sale of one product line and starting a different product line. These preliminary results suggest that local knowledge or personal contact may be important to induce changes in detailed practices, but new “big ideas” may be generated by more distance interactions. Third, Abebe, Fafchamps et al. (2018) report on an experiment in ‘learning by managing’ – where aspiring entrepreneurs spend four weeks working alongside management in successful medium-to-large firms. This model is inspired by the Hayekian notion that entrepreneurial knowledge – like riding a bike – might be

something that needs to be experienced in order to be learned. The authors find that their treatment has a large and significant effect on management quality, particularly concentrated upon incumbent firms.

4.3 Entrepreneurship and peer support

If entrepreneurial knowledge is easier to learn than to teach, then it may follow that successful peers – rather than trainers or consultants – are a more effective source of business wisdom. This is the basic motivation behind experiments testing for diffusion of management practices through random peer assignment and through mentoring.

Fafchamps and Quinn (2016) run an experiment in which managers of small enterprises are randomly assigned to judging committees for a business plan competition. The experiment, the first to exogenously create new social connections between firms' managers, was run with 345 judge/managers (average firm size six employees) in 2011 in Ethiopia, Tanzania and Zambia. By creating random links between managers, Fafchamps and Quinn were able to test directly the effects of new peers on firm outcomes. The only robust evidence in this regard was for VAT registration: being paired with a new peer who was VAT-registered at baseline increased the probability of VAT registration at follow-up by about seven percentage points, on a baseline mean of eight percent.¹⁴ The authors note that all three of the study countries were undergoing VAT reforms at the time of the experiment, increasing the salience of this issue. The general message from the experiment is that, at least for small and established firms, we should not expect substantial improvements in business practices merely through the creation of new network links.

Cai and Szeidl (2018) conducted a randomized experiment in China with a sample much larger and much more heterogeneous in firm size. The 2820 firm owners (of whom 1500 were treated) averaged 36 workers, and were all established within three years of the experiment's start in 2013. The treatment was fairly intensive; managers met monthly for a year. The 1500 treatment firms were organized into 150 groups of 10 firms each. The groups were of four types: (i) small, same sector, (ii) large, same sector; (iii) mixed size, same sector; and (iv) mixed size, mixed sector. An initial meeting was held with each group, after which the owners were encouraged to self-organize additional monthly meetings for the following year. Compliance with the request to organize and participate in these meetings was very high (87%), suggesting that owners found the meetings useful. The relatively large sample size and stratified design allows for quite a nuanced analysis of the results of the experiment. First, comparing the performance of enterprises in the treatment group with those in the control group, Cai and Szeidl find that the meetings lead to an increase in both the scale and profitability of the enterprises. Sales of treatment firms increase by 8 to 10 percent relative to the control firms, with comparable increases in material inputs, employment and assets. Why did the interactions lead to an increase in firm growth? Cai and Szeidl show evidence on several channels. First, there is evidence that firms share information on trading partners, with the number of referrals to trading partners and the number of direct relationships between firms in the group both significantly higher with treatment. Second, there is evidence the owners learn better management practices from one another. Management practices scores are 0.2 standard deviations higher in the treatment group compared with control.

¹⁴ The authors find significant diffusion of having a bank current account; this result is robust to corrections for multiple hypothesis testing within its family, but not across all outcomes tested.

Evidence that the outcomes were the result of learning from other firms in the group comes from two sources. First, firms randomized into groups with higher-quality peers showed larger increases in sales, profits, and management practices. Second, Cai and Szeidl carry out an additional experiment by providing selected members of each group information about either a savings product or a grant program. This additional experiment allows them to trace information flows more directly, by asking group members not receiving the information directly whether they are aware of the programs. The results on this additional experiment provide important lessons for scaling up the experiment: on average, information on the savings product flows through all groups and group members – but information on the grants flows only when the members are not direct competitors. The authors conclude with a cost-benefit analysis – showing that, even with conservative estimates, their experiment provides a cost-effective policy intervention.

Mentors and role models may be more effective than formal training because they provide localized information that is customized to the individual business. Brooks et al. (2018) design an experiment to test this proposition. Brooks et al.'s design compares standard business training with mentoring in a sample of 372 female-owned microenterprises on the outskirts of Nairobi, Kenya. The sample is split into a group receiving classroom training (which shows no effect), a group assigned to mentors and a control group. The mentors are successful business owners from the same community and sector of activity as the treated enterprise. The mentor-mentee pairs were encouraged to meet weekly at the mentor's place of business four times over a month, though many pairs continued to meet for more than a year beyond the official treatment period. The experiment shows a substantial short-term effect of mentoring, with profits of treated enterprises increasing by 20 percent. But the effects disappear about a year after the treatment begins. The mentor treatment appears to be effective while the mentor-mentee relationship is active and disappears as the incentives provided to the mentors are removed. More generally, a related literature tests for peer effects coming directly through the process of business training. For example, Field et al (2016) study a business counseling program in India, randomly varying whether participants were invited to attend with a friend; they find significant effects on business activity, but only for those who were trained with a friend.¹⁵

4.4 Demand-led growth

Our discussion so far has focused mainly on supply side factors. This is in keeping with the literature, which largely takes a supply-side approach to these issues. But it has long been recognized that shocks to demand can stimulate growth of firms and job creation generally.¹⁶ More recently, several studies have attempted to isolate the effect of demand shocks on firm productivity and growth.

Demand is limited by size of the market. For example, if transportation across two markets is prohibitively expensive, then a high productivity producer in Market A will not capture customers in Market B. Ghani et al. (2016), for example, show that the opening of the Golden Quadrilateral highway system in India led to a relocation of production to districts closer to

¹⁵ See also LaFortune, Peticar and Tessada (2018), who test for peer effects in a labor market training program in Chile.

¹⁶ See, for example, Tandler (1996, 1997).

the highway. By lowering input costs and increasing access to markets, the highway system shifted output to more productive firms.

But transportation costs are not the only reason that high-productivity producers find it difficult to capture customers from lower-productivity producers. Customers may lack trust in unfamiliar suppliers (McMillan and Woodruff, 1999), or lack the information about prices or quality necessary to make informed decisions. Andrabi et al. (2017) conduct an experiment in private schools in Pakistan, providing parents with “report cards” on achievements of students in the schools. Improving the information available to parents leads private schools to increase quality and/or lower prices, suggesting a more competitive market. Jensen and Miller (2018) use a natural experiment that increased the spread of information about the price and quality of boats purchased by fishermen in Kerala. Prior to wide adoption of mobile phones, fishermen both sold their catch and purchased their boats almost exclusively in local markets. Jensen (2007) shows that mobile phones allowed fishermen to gather information on prices being paid in various markets along the coast before landing with their catch. This led them to sell more frequently in markets outside their own village. Jensen and Miller (2018) show that the interaction with fishermen in other villages increased the knowledge on the quality and prices of boats built in those villages. This result was increased demand for the most productive boat builders and a loss of demand, and often exit, for the least productive boat builders. They estimate that productivity in the boat building industry increased by more than a quarter.¹⁷

Government is the largest purchaser of goods in almost every economy. Ferraz, Finan and Szerman (2015) use data from the national procurement system in Brazil to examine the effect on firm growth of demand shocks created by winning procurement contracts. Brazil holds an online auction for procurement of goods by the national government. In a quasi-experiment, the authors compare post-auction trajectories of firms that win and the second-highest bidder in a subset of auctions where the price gap between the bids is less than 1 percent and both bids were received with 30 seconds of the (random) ending time for the auction. They are able to link the procurement outcomes to matched employer-employee data, which allows them to observe firm growth subsequent to winning or just losing the auction. They find that winning an auction is associated with a modest but significant increase in employment (2.2 percentage points) an effect that, importantly, endures well beyond the life of the specific contract.

Demand shocks may come from access to foreign as well and domestic markets.¹⁸ Atkin et al. (2017) conduct an experiment aimed at opening international markets to a randomly selected subset of very small carpet-makers in Egypt.¹⁹ The researchers work through an NGO that attempts to link artisanal producers to buyers in high-income countries. The sample firms made carpets for the domestic market, and were offered orders to produce a very different style of carpet for customers in high-income countries. Atkin et al. show that linking to the

¹⁷ Interestingly, as the most productive firms grew in size, their workers became more specialized as well, working on fewer tasks in the production process.

¹⁸ There is an extensive non-experimental literature on the importance of access to foreign markets on the dynamics of firms growth. The standard prediction is in line with the Jensen and Miller results: access to foreign markets allows the most productive firms to grow while the least productive shrink and exit. De Loecker and Goldberg (2014) provide a review.

¹⁹ The sample has 303 firms; on average, these firms have spent 37 years in the carpet business, and, on average, have just one permanent employee each.

export market led firms to increase product quality and revenue per hour, even though productivity measured as (quality unadjusted) square feet per hour decreased. Most impressively, the increase in quality came without any formal training, and without any significant investment in capital stock. Rather, the authors show diffusion of techniques from the intermediary. Specifically, the authors can show that “treatment firms improve quality most along particular quality dimensions discussed during meetings between the intermediary and the producer”. While the experiment generates unique and quite rich insight into the underlying changes from the combined shock to demand and quality, the direct implications for development policy are less clear. As the authors note, the NGO indicates that six in seven attempts to build export links fail and, indeed, this case would likely have been one of the six if the researchers had not been persisting partners.

4.5 Firms and markets: The firm frontier

As the previous section illustrates, there is increasing evidence that demand is an important determinant of small firm dynamics. Frictions in goods markets arise from many sources – inadequate transportation networks, weak information flows, and lack of formal contract enforcement undermining trust in unknown trading partners. These frictions weaken the power of the market to reward the most productive firms, and to punish the least productive. This insight presents a promising new frontier for field experiments with firms: experiments that vary conditions at the level of local markets can quantify key aspects of the relationship between firms and their competitors – and, in doing so, may generate new insights on how policy can leverage market forces to improve firm productivity.

This remains a new area of experimental research, and one that is yet to yield clear and generalizable conclusions. Nonetheless, some results from initial experiments are worth noting. Initial work shows that local competition can matter for firm performance – yet already it is clear the effect of competition will depend upon specific context. For example, as noted earlier, Andrabi et al. (2017) show that exogenous information shocks in local village markets in Pakistan increase competitive pressures on firms (in this case, schools), leading firms (schools) either to improve quality or lower prices. Busso and Galliani (forthcoming) find similar effects from randomizing entry of small retail shops at the market level in the Dominican Republic. But increased competition may not always lead to similar improvements: Bergquist (2017) uses a field experiment to elicit measures of pass-through and demand curvature in maize markets in rural Kenya, showing that increasing the number of sellers in markets has little effect on prices and that sellers pass through only a small share of cost reductions.²⁰ Exploring more generally the price effects of exogenous quantity shocks in locally isolated markets, Cunha et al (forthcoming) exploit random variation between cash- and in-kind transfers in Mexico. They show that in-kind transfers cause a fall in prices, and find that cash transfers lead to higher prices only in more remote villages. While their data do not allow them to sort out the exact channels, these results are consistent with frictions between local sellers and their suppliers. These initial results highlight how little we understand about the process of price-setting in markets. (Given the number of factors that

²⁰ Casaburi and Reed (2017) conduct a pass-through experiment similar to one of Berquist’s experiments, and find that sellers increase the provision of credit but do not pass through the lower costs. Mitra et al (2018) provide price information to potato farmers in a random subset of villages in West Bengal and show that pass through increases as a result, consistent with a market characterized by bargaining between farmers and middlemen.

can differ across markets – information, social embeddedness, the importance of quality, etc. – we are likely to find that the shocks generated by experiments reveal a range of outcomes

The complexity of market interactions also explains why firms themselves struggle to optimize. Karlan and Zinman (forthcoming) work with a large microlender in Mexico (Compartamos) to randomize interest rates at the market level. Using data from the lender’s own administrative records, they show that credit demand elasticities are larger than expected – especially in the longer-run – and, indeed, are too large to be consistent with profit maximization. But they are also able to use data from credit bureaus covering all loans made in the market to understand how the price change by one lender affects behavior of other lenders. They find no crowding out of lending. Moreover, drawing on data from a ‘mystery shopping’ exercise, they show that competing lenders do not reduce rates in response to the larger discount.²¹ The authors argue that this is due to information frictions, such that (among other effects) “competitors face substantial risks and information gaps in evaluating whether it would be valuable to change prices”.

Together, these results raise important empirical questions about how competitive forces vary across different markets – questions that relate closely to traditional issues in the study of industrial organization. The increasing availability of administrative data opens exciting opportunities for designing experiments with individual firms that illuminate market-level interactions, using either reduced-form and structural approaches.

5. Conclusion

We have demonstrated the value of experiments with firms by reviewing how experiments have furthered our understanding of constraints to firm growth in developing countries. By increasing our understanding of how entrepreneurs respond to changes in circumstances, experiments provide us knowledge at a level of detail that is conducive to informing policy. The experiments put flesh on the bones of the models showing the magnitude of productivity dispersion in lower-income countries.

Of course, experiments rarely provide a full answer to policy questions, either; indeed, even in a review focused on experiments, we have highlighted some non-experimental work. Experiments are more feasible in some contexts (capital in very small firms, for example) than in others (the importance of capital in very large firms, for example). Moreover, most often experiments will uncover partial equilibrium outcomes. In this context, they are most illuminating when they uncover patterns in individual decision-making, showing how agents respond to changes in circumstances or incentives.

The power of experiments with firms is well illustrated by the trajectory of research on microenterprise training. More than a dozen rigorous evaluations of training programs were reviewed by McKenzie and Woodruff (2013). Any one of these experiments in isolation had limited, context-specific, implications. But collectively they showed that the mainstream

²¹For other examples of creative use of administrative data to understand market-level behavior among firms in developing countries, see Bernstein et al (2018) (who use fluctuations in global commodity prices as a natural experiment for the creation of new firms in Brazil), and Higgins (2018) (who exploits the exogenous roll-out of debit cards in Mexico to test effects on small retailers’ adoption of point-of-sale terminals).

training programs were not effective. This led researchers and practitioners alike to reconsider how training should be delivered, and to innovate. Recent work shows that at least some of these innovations are promising. Of course, positive results in a single specific context should carry no more weight than negative results in a single specific context. Policy formulation needs replications, which are often under-delivered given the incentives of researchers.

Experiments are most useful when they remain tethered to theory, for at least two reasons. First, theory has an important role in suggesting new experimental designs. For example, as the previous section illustrated, recent experiments have provided important new insights into the way that demand can affect firm growth – but nothing from experiments on the supply side pointed researchers to this issue of demand. Second, while results from individual experiments are always context-specific, theory can help us to understand how and when experimental results are likely to generalize. For this reason, we view the recent movement towards combining experiments with structural models as promising. It is an important step in addressing concerns with external validity and, in some contexts, to addressing general equilibrium effects. The combination of theory and experimentation has the potential to provide a connection between the coherent big-picture view of the macro-development literature and the mosaic created by individual experiments.

Bibliography

Abebe, G, Caria, S, Ortiz-Ospina, E. 2017. The Selection of Talent: Experimental and Structural Evidence from Ethiopia. Working paper.

Abebe G, Caria, S, Fafchamps, M, Falco, P, Franklin, S, Quinn, S. 2018. Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City. Working paper.

Abebe, G, Fafchamps, M, Koelle, M, Quinn, S. 2018. Learning by Managing: A Field Experiment with Aspiring Entrepreneurs. Working paper.

Acemoglu D, Johnson S, Mitton T. 2009. Determinants of Vertical Integration: Financial Development and Contracting Costs. *Journal of Finance*, 64(3):1251-1290

Ackerberg DA, Caves K, Frazer G. 2015. Identification properties of recent production function estimators. *Econometrica*, 83:2411–2451.

Alfonsi L, Bandiera O, Bassi V, Burgess R, Rasul I, Sulaiman M, Vitali A. 2017. Tackling youth unemployment: evidence from a labor market experiment in Uganda. Suntory and Toyota International Centres for Economics and Related Disciplines, London School of Economics.

Anderson-MacDonald SJ, Chandy R, Zia B. 2017. Pathways to profits: the impact of marketing vs. finance skills on business performance. *Management Science*, forthcoming.

Andrabi, Tahir, Jishnu Das and Asim Ijaz Khwaja. 2017. "Report Cards: The Impact of Providing School and Child Test Scores on Educational Markets. *American Economic Review*, 107(6): 1535-1563

Angelucci M, Karlan D, Zinman J. 2015. Microcredit impacts: Evidence from a randomized microcredit program placement experiment by Compartamos Banco. *American Economic Journal: Applied Economics*, 7(1):151-82.

Arellano M, Bond, S. 1991. Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. *Review of Economic Studies*, 58(2): 277-97.

Atkin D, Khandelwal AK, Osman A. 2017. Exporting and Firm Performance: Evidence from a Randomized Experiment. *The Quarterly Journal of Economics*, 132(2):551–615.

Atanasio O, Augsburg B, De Haas R, Fitzsimons E, Harmgart H. 2015. The impacts of microfinance: Evidence from joint-liability lending in Mongolia. *American Economic Journal: Applied Economics*, 7(1):90-122.

Augsburg B, De Haas R, Harmgart H, Meghir C. 2015. The impacts of microcredit: evidence from Bosnia and Herzegovina. *American Economic Journal: Applied Economics*, 7(1):183-203.

Bandiera O, Hansen S, Pratt A, Sadun R. 2017. CEO behavior and firm performance. Harvard Business School Working Paper 17-083.

Banerjee A, Duflo E. 2014. Do firms want to borrow more? Testing credit constraints using a directed lending program. *The Review of Economic Studies*, 81(2):572–607.

Banerjee A, Duflo E, Glennerster R, Kinnan C. 2015. The miracle of microfinance? Evidence from a randomized evaluation. *American Economic Journal: Applied Economics*, 7(1):22-53.

Banerjee A, Karlan D, Zinman J. 2015. Six randomized evaluations of microcredit: Introduction and further steps. *American Economic Journal: Applied Economics*, 7(1):1-21.

Beaman, L. and J. Magruder (2012). Who Gets the Job Referral? Evidence from a Social Networks Experiment. *The American Economic Review* 102(7), 3574–3593.

Beaman, L., N. Keleher, and J. Magruder (2018). Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi. *Journal of Labor Economics* 36(1), 121–157.

Bergquist, L.F. 2017. Pass-through, Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya. Working paper.

Bernstein, S., Colonnelli, E., Malacrino, D., and McQuade, T. 2018. Who Creates New Firms When Local Opportunities Arise? *NBER Working Paper No. 25112*.

Blattman C, Dercon S. 2018. The impacts of industrial and entrepreneurial work on income and health: Experimental evidence from Ethiopia. *American Economic Journal: Applied Economics*, 10(3):1-38.

Bloom N, Sadun R, Van Reenen J. 2012. Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review*, 102(1):167-201.

Bloom N, Van Reenen J. 2007. Measuring and Explaining Management Practices Across Firms and Countries. *The Quarterly Journal of Economics*, 122(4):1351–1408

Boehm L, Oberfield E. 2018. Misallocation in the market for inputs: enforcement and the organization of production”. Working paper, Princeton.

Boettke PJ. 2002. Information and Knowledge: Austrian Economics in Search of its Uniqueness. *The Review of Austrian Economics*, 15:263

Blundell R, and Bond S. 1998. Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*. 87: 115-143.

Brooks W, Donovan K, Johnson TR. 2018. Mentors or teachers? Microenterprise training in Kenya. *American Economic Journal: Applied Economics*. In press.

Bruhn M, Karlan D, Schoar A. 2018. The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in Mexico. *Journal of Political Economy*, 126(2):635-687.

Busso M, and Galiani, S. Forthcoming. The Causal Effect of Competition on Prices and Quality: Evidence from a Field Experiment. *American Economic Journal: Applied*.

Cai J, Szeidl A. 2018. Interfirm relationships and business performance. *The Quarterly Journal of Economics*, 133(3):1229-1282.

Campos F, Frese M, Goldstein M, Iacovone L, Johnson HC, McKenzie D, Mensmann M. 2017. Teaching personal initiative beats traditional training in boosting small business in West Africa. *Science*, 357(6357):1287-1290.

Casaburi, L., and Reed, T. 2017. Competition in Agricultural Markets: An Experimental Approach. Working paper.

Crepon B, Devoto F, Duflo E, Parienté W. 2015. Estimating the impact of microcredit on those who take it up: Evidence from a randomized experiment in Morocco. *American Economic Journal: Applied Economics*, 7(1):123-50.

Dalton P, Bilal Z, Rüschenpöhler J, Uras B. 2018. Learning business practices from peers: experimental evidence from small-scale retailers in an emerging market. DFID Working Paper.

De Loecker J, Goldberg PK. 2014. Firm Performance in the Global Market. *Annual Review of Economics*, 6:201–227.

De Mel S, McKenzie D, Woodruff C. 2008. Returns to capital in microenterprises: evidence from a field experiment. *The Quarterly Journal of Economics*, 123(4):1329–1372.

De Mel S, McKenzie D, Woodruff C. 2012. One-time transfers of cash or capital have long-lasting effects on microenterprises in Sri Lanka. *Science*, 335:962–966.

De Mel S, McKenzie D, Woodruff C. 2013. The Demand for, and Consequences of, Formalization among Informal Firms in Sri Lanka. *American Economic Journal: Applied Economics*, 5(2):122-150,

De Mel S, McKenzie D, Woodruff C. 2018a. Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises. *American Economic Journal: Applied Economics*, forthcoming.

De Mel S, McKenzie D, Woodruff C. 2018b. Micro-equity for Microenterprises. Working paper, Oxford.

Dlugosch TJ, Klinger B, Frese M, Klehe UC. 2018. Personality-based selection of entrepreneurial borrowers to reduce credit risk: Two studies on prediction models in low- and high-stakes settings in developing countries. *Journal of Organizational Behavior*, 39(5):612–628.

Fafchamps M, McKenzie D, Quinn S, Woodruff C. 2014. Microenterprise growth and the flypaper effect: Evidence from a randomized experiment in Ghana. *Journal of Development Economics*, 106:211-226.

Fafchamps M, Quinn S. 2016. Networks and manufacturing firms in Africa: Results from a randomized field experiment. *World Bank Economic Review*.

Fafchamps M, Woodruff C. 2017. Identifying gazelles: expert panels vs. surveys as a means to identify firms with rapid growth potential. *World Bank Economic Review*, World Bank Group, 31(3):670-686.

Ferraz C, Finan F, Szerman D. 2015. Procuring firm growth: the effects of government purchases on firm dynamics. NBER Working Papers 21219.

Field E, Jayachandran S, Pande R and Nigol R. 2016. Friendship at Work: Can peer effects catalyze female entrepreneurship? *American Economic Journal: Policy* 8(2): 125-153.

Field E, Pande R, Papp J, Rigol N. 2013. Does the classic microfinance model discourage entrepreneurship among the poor? Experimental evidence from India. *American Economic Review* 103(6):2196-2226.

Fischer G. 2013 Contract structure, risk-sharing, and investment choice. *Econometrica* 81:883-939.

Gandhi A, Navarro S, Rivers D. 2016. On the Identification of Gross Output Production Functions. *Journal of Political Economy*, forthcoming.

- Ghani E, Goswami AG, Kerr WR. 2016. Highway to success: the impact of the golden quadrilateral project for the location and performance of Indian manufacturing. *The Economic Journal*, 126:317-357.
- Gibbons R, Henderson R. 2012. Relational Contracts and Organizational Capabilities. *Organization Science*, 23(5):1350-1364.
- Hardy M, McCasland J. 2017. Are small firms labor constrained? Experimental evidence from Ghana. Working paper. NYU-Abu Dhabi.
- Hayek FA. 1945. The use of knowledge in society. *The American economic review*, 35(4):519-530.
- Hayek FA. 1988. *The Fatal Conceit: The Errors of Socialism*. W. W. Bartley III (ed.), Chicago: The University of Chicago Press.
- Higgins, S. 2018. Financial Technology Adoption. Working paper.
- Hsieh CT, Klenow PJ. 2009. Misallocation and manufacturing TFP in China and India. *The Quarterly journal of economics*, 124(4):1403-1448.
- Hsieh CT, Olken BA. 2014. The Missing "Missing Middle". *Journal of Economic Perspectives*, 28(3):89-108.
- Hussam R, Rigol N, Roth BN. 2017. Targeting high ability entrepreneurs using community information: mechanism design in the field. PEDL Working Paper.
- Jensen R. 2007. "The Digital Divide: Information (Technology), Market Performance and Welfare in the South Indian Fisheries Sector." *Quarterly Journal of Economics*, 122(3): 879-924.
- Jensen R, Miller NH. 2018. Market Integration, Demand and the Growth of Firms: Evidence from a Natural Experiment in India. *American Economic Review*, forthcoming.
- Kahneman D, Klein G. 2009. Conditions for intuitive expertise: a failure to disagree. *American Psychologist*, 64:515-526.
- Karlan D, Zinman J. 2011. Microcredit in theory and practice: using randomized credit scoring for impact evaluation. *Science*, 332(6035):1278-1284.
- Karlan D, Zinman J. Forthcoming. Long-run price elasticities of demand for microcredit: Evidence from a countrywide field experiment in Mexico. *Review of Economic Studies*.
- Kerr W, Lerner J, Schoar A. 2014. The consequences of entrepreneurial finance: evidence from angel financings. *The Review of Financial Studies*, 27(1):20-55.
- Kerr, W, Nanada, R, Rhodes-Kropf, M. 2014. Entrepreneurship as Experimentation. *Journal of Economic Perspectives*. 28(3): 25-48.

- Klinger B, Schündeln M. 2011. Can entrepreneurial activity be taught? Quasi-experimental evidence from Central America. *World Development*, 39(9):1592-1610.
- Lafortune J, Perticarà M and Tessada, J. 2018. The Benefits of Diversity: Peer effects in an adult training program in Chile. *Working paper.*
- Lagakos D. 2016. Explaining cross-country productivity differences in retail trade. *Journal of Political Economy*, 124(2):579-620.
- Lagakos D, Moll B, Porzio T, Qian N, Schoellman T. 2018. Life cycle wage growth across countries. *Journal of Political Economy* 126(2):797-849.
- La Porta R, Lopez-de-Silanes F, Shleifer A, Vishny R. 1997. Trust in large organizations. *The American Economic Review*, 87(2):333-338.
- Leaven L, Woodruff C. 2007. The quality of the legal system, firm ownership, and firm size. *Review of Economics and Statistics*, 89(4):601-614.
- Lemos R, Scur D. 2018. All in the family? CEO choice and firm organization. CEP Discussion Papers, CEPDP1528. Centre for Economic Performance, London School of Economics and Political Science, London, UK.
- Levinsohn J, Petrin A. 2003. Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, 70(2):317-341.
- Macchiavello R, Menzel A, Rabbani, A. and Woodruff, C. 2018. Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector. Working paper, Oxford University.
- McCaig, B, Pavcnik, N. Out with the Old and Unproductive, in with the New and Similarly Unproductive. GLM LIC Working Paper 23.
- McMillan J, Woodruff C. (1999). Interfirm relationships and informal credit in Vietnam. *The Quarterly Journal of Economics*, 114(4), 1285-1320.
- McKenzie D. 2017a. How effective are active labor market policies in developing countries?: A critical review of recent evidence. *The World Bank Research Observer* 32(2):127-154.
- McKenzie D. 2017b. Identifying and spurring high-growth entrepreneurship: experimental evidence from a business plan competition. *American Economic Review*, 107(8):2278-2307.
- McKenzie D, Sansone D. 2017. Man vs. machine in predicting successful entrepreneurs: evidence from a business plan competition in Nigeria”. Policy Research Working Paper no. WPS 8271.
- McKenzie D, Woodruff C. 2008. Experimental evidence on returns to capital and access to finance in Mexico. *The World Bank economic review* 22(3):457-482.

McKenzie D, Woodruff C. 2013. What are we learning from business training and entrepreneurship evaluations around the developing world? *The World Bank Research Observer* 29(1):48–82.

McKenzie D, Woodruff C. 2017. Business practices in small firms in developing countries. *Management Science*, 63(9):2967-2981.

Meager R. 2018a. Aggregating distributional treatment effects: A Bayesian Hierarchical Analysis of the Microcredit Literature. Working paper, London School of Economics

Meager R. 2018b. Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments. *American Economic Journal – Applied*, forthcoming.

Mitra, S. Mookherjee, D. Torero, M. ad Visaria, S. 2018. Asymmetric Information and Middleman Margins: An Experiment with Indian Potato Farmers. *The Review of Economics and Statistics*, 100(1), pp. 1–13

Oğuz F. 2010. Hayek on tacit knowledge. *Journal of Institutional Economics*, 6(2):145–165.

Olley G, Pakes A. 1992. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263-1297.

Rajan RG, Zingales L. 1998. Financial dependence and growth. *American Economic Review* 88(3):559-586.

Tarozzi A, Desai J, Johnson K. 2015. The impacts of microcredit: Evidence from Ethiopia. *American Economic Journal: Applied Economics*, 7(1):54-89.

Tendler J. 1997. *Good government in the tropics* (1st. pbk. ed). Johns Hopkins University Press, Baltimore.

Tendler J, Alves M. 1996. Small firms and their helpers: lessons on demand. *World Development*, 24(3):407-426.

Ulysea, G. 2018. Firms, informality, and development: theory and evidence from Brazil. *American Economic Review*, 108(8):2015-47.

Witte, M. 2018. “Job Referrals and Strategic Network Formation – Experimental Evidence from Urban Neighbourhoods in Ethiopia,” Working paper, Oxford.

Yunus M. 1989. Grameen Bank: organization and operation. In: Jacob Levitsky (ed.) *Microenterprises in Developing Countries*. London: Intermediate Technology Publications.