From Micro to Macro in an Equilibrium Diffusion Model*

Wyatt Brooks
Arizona State University

Kevin Donovan
Yale University

Terence R. Johnson
University of Virginia

January 2023

Abstract

We quantify the benefits of better firm-to-firm matching in an aggregate diffusion model and use it to interpret empirical moments from interventions that do so at a smaller scale. We build a model in which individuals reap profitable knowledge from others in the economy, then study the implications when meeting technology makes it easier to meet high-knowledge agents. We estimate the model using recent empirical evaluations of small-scale programs that create new learning opportunities among firm managers. The equilibrium gains from better meetings are primarily disciplined by empirical moments other than the average treatment effect. Extrapolating from it therefore generates large quantitative bias. However, for a broad class of diffusion models and interventions, these additional moments can be estimated with a simple linear regression from experimental data, thus providing practical information on the potential at-scale gains that can be useful for policy decisions.

JEL Classification Codes: O11, O33, O40
Keywords: diffusion, general equilibrium, ability

*This paper has benefited from comments and discussions by participants at the Bank Al-Maghrib Conference on Small Firm Growth, the IESR Conference on Firms in Emerging Economies, the NBER Growth Meeting, the Society for Economic Dynamics, Y-RISE, and various seminars, especially those from Daron Acemoglu, Paco Buera, Lorenzo Caliendo, Teresa Fort, Mushfiq Mobarak, Ezra Oberfield, Jesse Perla, Peter Schott, Kevin Song, Meredith Startz, Mike Waugh, and Daniel Xu. We especially thank Jonas Heipertz, Melanie Morten, and Tommaso Porzio for insightful discussions. Brian Mukhaya provided outstanding research assistance. *Contact Information:* Brooks, wjbrooks@asu.edu; Donovan, kevin.donovan@yale.edu; Johnson, trj2j@virginia.edu

1 Introduction

One of the fundamental constraints to economic development is limited managerial capacity. Despite the ubiquity of small firms in developing countries, many have low profit and few workers. Billions of dollars are spent attempting to correct these skill deficiencies (Blattman and Ralston, 2017). While many solutions have been proposed, a flurry of recent micro-level interventions have highlighted one promising channel: learning from others. These interventions enable individual business owners or managers to interact with people that they may not have had the chance to otherwise, finding increases in profit, technology adoption, and managerial skills.[1]

Yet the general equilibrium, at-scale implications of these interventions are less explored. There is reason to suspect they may differ from what could be extrapolated directly from these interventions. When the knowledge created in these meetings can further permeate the economy through the diffusion, the gains from better matching will depend on changes to the entire distribution of knowledge (Lucas and Moll, 2014; Perla and Tonetti, 2014). These changes need not be directly related to the gains to the average person in the intervention. Our contribution is to investigate these relationships in a general equilibrium model of firm-to-firm diffusion, in which we study the cost of frictions that limit learning opportunities and how they relate to these micro-level interventions.

We focus our model around the results of a randomized controlled trial in Kenya from earlier work (Brooks et al., 2018), though as we will discuss at length later, many of the insights can be formally generalized to a class of models and interventions that cover a substantial amount of recent work. In this RCT, we randomly paired high- and low-profit female business owners in Nairobi, Kenya. On average, profit rose by 19 percent for the less-profitable member of the match, with no statistically significant change for the more profitable business owner. The changes mostly occurred on the cost side: treated owners were more likely to switch suppliers and their costs declined significantly.

Our RCT, and many like it, focused on outcomes of those engaged in a match. This has the benefit of clearly showing that there exists some friction in the learning process that can potentially be overcome with policy. But the total gains of such a program when implemented at scale should include the fact that these agents learn

[1] There now exists a broad set of interventions that are designed to give firms access to new information embedded in other economic agents. Cai and Szeidl (2018) create randomly-formed business groups in China. Brooks et al. (2018) and Lafortune et al. (2018) create random one-to-one matches among Kenyan and Chilean microenterprise owners. Fafchamps and Quinn (2018) introduce smaller-scale firm owners to managers of high-profit firms in Ethiopia, Tanzania, and Zambia. Relatedly, Atkin et al. (2017) introduce random variation in buyer-seller links in Egypt and Beaman et al. (2021) randomly seed information on new technology in Malawi. See also Giorcelli (2019) and Bianchi and Giorcelli (2022) for historical examples of firm-to-firm knowledge transfers and Munshi (2008) and Breza et al. (2019) for reviews.

new information and skills that can then be transmitted to others in the future. We formalize this equilibrium diffusion by building a general equilibrium model of knowledge diffusion. The basic features of the model share much with standard models in the literature, such as those reviewed in Alvarez et al. (2008) and Buera and Lucas (2018). We adopt the specifics to take seriously the channels highlighted by the RCT and the economic realities of the economy in which it is run.

Specifically, we build a model in which individuals can choose to run a firm or work at one. This decision depends on the income that can be earned in either occupation. Self-employment income depends on knowledge of the market. More knowledgeable agents can more easily seek out low-cost suppliers, consistent with our empirics. We refer to this as "knowledge" to fix definitions, but it is more broadly any ability or skill that makes the firm more profitable. This knowledge can be diffused. Each period, every agent randomly interacts with a firm owner. If she is lucky enough to meet a high-knowledge firm owner, she can increase her own stock of knowledge, thus allowing her to more easily find lower cost suppliers. Thus, the distribution of knowledge in the economy is the key aggregate state variable. It controls the potential set of meetings available to agents, and is determined both by the economic decisions made by agents and the technology that translates meetings into knowledge.

We discipline the model in large part using evidence from the RCT. We match a number of relevant empirical moments from the RCT, including the average treatment effect and measures of heterogeneity within the treatment effect. These pin down important features of the learning technology and do so in a way that is independent of much of the remaining model structure we impose to measure quantitative outcomes. The model also matches other moments not directly targeted like the fade-out of the treatment effect over time.

We use the model to simulate the gains from making it easier for firm owners to learn from each other, taking into account general equilibrium forces. These forces are two-fold. First, there is a direct effect of the better meeting technology. If everyone can more easily learn from others, this pushes out the equilibrium distribution of knowledge and makes all firms more profitable. The second equilibrium force is on prices. By making firms more profitable, the wage increases. This competitive pressure pushes marginal firms out of business and amplifies the direct effect. Pushing out the lowest-knowledge firms de-congests learning for everyone else. Quantitatively, we find that average income rises by 11 percent; two-thirds comes from the direct effect on the distribution of knowledge and the remainder coming through the amplification effect from price changes.

We then use the model to ask how much discipline the empirical average treatment effect puts on the realized general equilibrium gains from the model. This moment is often used to measure the success of an intervention or its potential for scale, but a growing literature has pointed out that the link to general equilibrium implications may not be straightforward (see reviews in Buera et al., 2022; Lagakos and Shu, 2022). To answer this question, we re-compute the aggregate gains, but drop all treatment effect moments used in estimation except the average treatment effect. Instead of using these additional moments, we arbitrarily set the parameters that are matched to them. When we do this, we find that the aggregate gains can vary from 0.6 to 38 percent depending on the parameters chosen. These results have two implications. First, our modeling choices do not pre-determine small or large aggregate gains. They can vary widely depending on which way the empirics push the model parameters. Second, the discipline on the aggregate gains comes in large part from moments other than the average treatment effect.

When we investigate these other empirical moments, we find that an adjusted covariance moment plays a critical role in generating the equilibrium gains. In the data, this moment represents how the individual-level treatment effect co-varies with the baseline profitability of individual's match. Or put differently, it measures how much farther into the tail of the knowledge distribution an agent is pushed by meeting with a high-knowledge match relative to a low-knowledge match. This empirical notion is closely connected to the driving force behind the equilibrium gains we observe in the model. Here we shock the meeting technology to make the average match better (as we do in the experiment) and the equilibrium gains are driven primarily by how much the knowledge distribution shifts in response, similar to what our covariance moment measures. Indeed, we show that changes to the model parameter pinned down by this moment cause larger shifts in the knowledge distribution than the parameter pinned down by the average effect.

As a practical matter, we show that this covariance moment can be measured as a coefficient from a simple linear regression run on treated firms within the experiment. Thus, one contribution of this paper is to highlight an important policy-relevant moment that can (1) be measured easily in a partial equilibrium RCT and (2) provides important information about the gains that could be achieved at scale.

Motivated by this last result, our final exercise is to take a step back and ask if there are any more general lesson to be drawn about measurement in partial equilibrium RCTs in the presence of diffusion. We view this as an important question given the growing set of RCTs that have a similar flavor to the one used here (i.e., random opportunities to learn) but differ in exact implementation (i.e., group meetings of firm managers as in Cai and Szeidl (2018) or buyer-supplier links in Atkin et al. (2017)). We show the same procedure we use holds for a broad class of both equilibrium diffusion models and RCTs. That is, our procedure leverages features of a broader class of interventions than the specifics of our Kenyan RCT, and pins down the same relationship between moments and parameters in a broader class of diffusion models than the one we use.

Therefore, while we join a recent literature that disciplines macro-development models with experiments to simulate scaled policy (e.g. Buera et al., 2021; Fujimoto et al., 2021; Lagakos et al., 2021; Kaboski et al., 2022; Fried and Lagakos, 2021), we further show how to measure additional moments *within* the RCT that provide relevant information about the gains at scale that are useful for policy decisions. These results show that "small scale" RCTs can provide relevant information about at-scale gains with careful measurement and links to theory, without necessarily incurring the costs required to run larger and larger clustered experiments to measure equilibrium effects.[2] We further link together a growing micro-development literature with a parallel one in macro-development and growth, where the non-rivalrous nature of information has long been seen a contributor to aggregate growth (Romer, 1990; Jones, 1995; Kortum, 1997). Jovanovic and Rob (1989), Lucas (2009), Lucas and Moll (2014), and Perla and Tonetti (2014) micro-found the equilibrium diffusion of knowledge between agents, and form the basis for the model we use here.[3]

2 Empirical Evidence on Diffusion

Recently, several microeconomic studies have documented the potential benefits for firm owners or managers who are randomly chosen to interact with highly skilled individuals. These take the shape of other managers or potentially by interacting through supply chains. We focus on one of those here and take up others in the Appendix. We utilize the randomized controlled trial in Brooks et al. (2018).[4] Space constraints naturally require us leave out some details, but we refer interested readers to Brooks et al. (2018) for other (less critical) details.

[2] In addition to costs, defining catchment areas for clusters is difficult. Muralidharan and Niehaus (2017) discuss the difficulty of defining spillover borders. Berguist et al. (2019) formalize a related issue when attempting to interpret agricultural interventions in the presence of trade flows.

[3] These models have been extended to international trade (Buera and Oberfield, 2020; Perla et al., 2021), innovation policy (Benhabib et al., 2021; Lashkari, 2020), and various types of learning among co-workers (Herkenhoff et al., 2018; Jarosch et al., 2021; Wallskog, 2021).

[4] We use this experiment because we have access to the relevant data, not because we view this particular experiment as better or worse than other similar experiments that fall under our assumptions.

2.1 Details of RCT and Data Collection

The experiment took place in Dandora, Kenya, a dense informal settlement on the outskirts of Nairobi. Via a representative sample of firm owners in Dandora, we randomly matched more profitable entrepreneurs with less profitable ones. We followed the owners over 17 months to measure changes in business practices and profitability over time.

The sample we draw from is the set of female business owners who have been in operation for less than 5 years.[5] We then randomly select a subset of these business owners to randomly match with a more profitable owner. A control group receives no such offer. Firms were then surveyed over 6 quarters to track the time series of treatment.

The set of more profitable owners were selected from those businesses with owners over 40 years old and at least 5 years of experience. This was to minimize the importance of "luck" in baseline profit realizations to allow us to focus on truly productive business owners. We then recruited business owners with the highest profit until we had a sufficient number for matches. Of those contacted, 95 percent accepted. Matches with the treatment firms were randomly created conditional on industry. Figure 7 in Appendix A shows the profit distributions for the full population, the sample of control and treatment firms, and the intervention-defined matches. As expected, our study population is somewhat poorer the average and the matches are drawn from the far right tail.

It will be useful when we relate these results to the model to have some relevant distributions defined. We refer to the cumulative distribution of baseline profit for treatment and control firms as $H_{T,\pi}(\pi)$ and $H_{C,\pi}(\pi)$ (which in theory are the same due to randomization). The baseline matches for the treatment group are distributed with c.d.f. $\widehat{H}_{T,\pi}(\hat{\pi})$.

Details of a "Match" Like any RCT, our matches were designed to remain faithful as possible to the model without sacrificing the practicality required to generate take-up. This necessarily involves trade-offs, which we discuss here. The program was pitched to both sides of the match as a mentee-mentor relationship, and thus was explicitly focused on business success. The more successful business owners were the "mentors," while the less successful were the "mentees." The mentors were told they could potentially help other business owners learn the requisite skills required to operate in Nairobi. We provided no topics to discuss, preferring that the content

[5] The sex selection criteria is to limit heterogeneity outside the model. Note, however, that females make up 65 percent of business owners in Dandora and 71 percent of owners with businesses open less than 5 years.

was self-directed. We offered mentors only some optional, vague prompts in an initial orientation meeting that could be used ("What challenges did your mentee face this week?"). Matches were designed to last for one month, though of course there was no restriction on meeting after the formal end of the program.

One potential concern is our meetings may not necessarily reflect those that underlie the usual matching process, perhaps related to indirectly priming mentees to believe the matches would be beneficial. There is little we can do to rule this out completely.[6] It is a problem common to most randomized controlled trials when attempting any extrapolation outside the study sample. We designed the details to be as light-touch as possible while offering a reasonable chance at successful take-up (i.e., to initiate matches, we simply gave the mentee the phone number of the mentor).

Results are balanced given the randomization and we provide details in Appendix A.

2.2 Treatment Effects and Underlying Mechanisms

Figure 1 begins by plotting the average treatment effect (as a percentage) over time, along with the 95 percent confidence interval. There is a large treatment effect in the immediate aftermath of the treatment period that fades quickly. We test whether our model rationalizes this pattern after laying out the model below, and find that it does.

Figure 1: Time Series of the Average Treatment Effect (from Brooks et al., 2018)

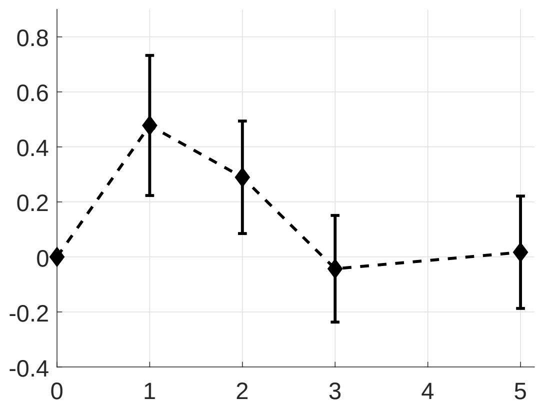

Figure notes: Figure plots average treatment effect as percentage above control mean (0.4 = 40%), along with the 95 percent confidence interval. Treatment takes place between quarters 0 and 1.

[6] We note, however, that evidence of the mentor's business success are easily visible to the mentee. Mentors had substantially more physical capital and workers, and had a fixed, relatively large buildings from which they conducted business. Moreover, the first meeting took place at the mentor's business. Thus, that the mentor was "good" at running a business would likely have been understood with or without us.

Mechanisms The observed changes in profit primarily come from lowering input costs, with unit inventory costs falling by 49 percent relative to the control group. Consistent with the cost channel, we find that treated firms are 19 percent more likely to switch suppliers in the aftermath of the treatment. We will take seriously this channel in the next section.

How does this result square with the quick fade-out observed in Figure 1? We find that the economic environment is characterized by substantial buyer-supplier turnover. Sixty-two percent of control firms switch suppliers in the 3 quarters immediately following treatment. Thus, any value procured by firms on this dimension by meeting with others is likely to have a short half-life. As we will show shortly, estimating the law of motion for ability similarly points to limited persistence, even without targeting this half-life directly.

Finally, we note that the matches create surplus and are not transfers between the two members in the match. We use the details of the matching procedure to estimate no changes in profitability, scale, or any management practices for the more productive member of the match.[7]

Is there diffusion? A natural remaining question is whether or not knowledge actually diffuses from this experiment. First, we note that changes in profit are driven by changes in managerial ability (Brooks et al., 2018). We observe no loans, joint input purchases or bulk discounts, profit sharing, or other mechanisms that would suggest alternative margins.

But an alternative theory of managerial capacity would be the high-knowledge match replacing the agent's span of control for the duration of the match. This would suggest a model of firm hierarchy (as in Caliendo et al., 2015), but perhaps not knowledge diffusion.

First, we note that our model can replicate the treatment effect pattern above (we discuss this Section 3.4), so a quick treatment effect fade-out is insufficient to answer this question. However, changes in ability could be driven by the more knowledgeable match replacing the agent's span of control for the duration of the match. This would suggest a model of firm hierarchy (as in Caliendo et al., 2015), but perhaps not knowledge diffusion.

First, we find that the change in profit is driven by changes in managerial ability (Brooks et al., 2018).[8] However, changes in ability could be driven by the mentor

[7] See Brooks et al. (2018) for details on the procedure. This forms the basis for our assumption of the max operator in the law of motion for ability, though as we discussed in Section 5.4, it is not crucial. Jarosch et al. (2021) find similar evidence among German co-workers.

[8] We also observe no loans, joint input purchases or bulk discounts, profit sharing, or other mechanisms that suggest alternative theories.

replacing the mentee's span of control for the duration of the match. This would suggest a model of firm hierarchy (as in Caliendo et al., 2015), but perhaps not ability diffusion.

In part to answer this question, we re-ran the field experiment used here with one additional feature: we randomly selected original mentees to mentor control firms *ex post*. We find that characteristics of the *original* mentor impact the treatment effect of the second-round mentee. This again points to the diffusion. These results form the basis of ongoing follow-up work to understand the details of how random meetings affect profitability, but we provide the relevant regression result in the Appendix B to provide evidence on this point.[9,10]

Together, this evidence shows that a diffusion model is the correct framework in which to study scaling these interventions.

3 A General Equilibrium Model of Diffusion

We now build a full general equilibrium model to study what we learn about the potential aggregate gains from this RCT. We build the model to remain consistent with the economic environment of the study and with the mechanisms highlighted in the previous section. Since the economy is primarily made up of relatively small firms, there is a non-trivial occupational choice decision between wage work and firm operation. We also model a search and bargaining process between firms and suppliers, motivated by the empirical results. Finally, we model a small open economy. Suppliers have access to infinite supply of inputs, consistent with their connection to the broader Nairobi and Kenyan economy. The labor market and diffusion of ability are local.

3.1 Economic Environment

Time is discrete and infinite. A period is one quarter. There is a unit mass of agents. The state of an agent is her ability to find a supplier z, which evolves over time. The aggregate state of the economy is the distribution of ability, $M(z)$. Each agent dies with exogenous probability δ and a mass δ of new agents replace them each period. New agents draw their initial ability from a fixed distribution with c.d.f. $G(z)$.

Every agent has flow utility $u(c, s) = \omega \log(c) + (1 - \omega) \log(1 - s)$, where c is consumption and s is effort (discussed below). ω is the relative weight of consumption

[9] This work is pre-registered with AEA Registry ID AEARCTR-0005564.

[10] In a similar RCT with randomized group meetings of Chinese managers, Cai and Szeidl (2018) cross-randomize information on a new financial product to test whether information is flowing between randomly matched groups of firms. They find that it is. They take this as direct evidence of transmission between firms, as do we. We estimate a model of this experiment in the Appendix.

in utility. There are no borrowing and savings markets, so consumption is equal to income.

Each period, an agent can choose between running a firm and working at one. Wage work earns the market-clearing wage w, as labor is traded on a competitive market. Firms earn profit $\pi = x^{\alpha} n^{\eta} - p_x x - wn$, where x and n are intermediate inputs and labor services, and p_x and w are their respective prices.

Supplier Matching, Input Choice, and the Importance of Ability Intermediate purchases require a firm to seek out a supplier. There are a continuum of suppliers indexed by their marginal cost m who earn profit $\pi^s(m,x) = (p_x - m)x$. Suppliers can source from some outside entity and thus can provide whatever amount of input x is requested by firms at marginal cost m. This is consistent with their connection to the broader Nairobi and Kenyan economy.

Firms exert search effort s to find a supplier. Given effort s and ability z, the firm matches with a supplier $m = e^{-s} z^{\frac{\alpha+\eta-1}{\alpha}}$. Thus, ability makes it easier to seek out lower cost suppliers.[11]

Once they meet, the two parties Nash bargain over the price p_x that the firm will pay, taking into account the optimal input choice by the firm at that price. This implies an equilibrium price

$$p_x^*(m) = \text{argmax}_{p_x} \quad (\pi)^{\nu} (\pi^s)^{1-\nu} \tag{3.1}$$

for bargaining weight ν and the participation constraint that both the supplier and firm earn weakly positive profit.

Ability Transmission Ability is transmitted between agents. It evolves as

$$z' = e^{c+\varepsilon} z^{\rho} \max\left\{1, \frac{\hat{z}}{z}\right\}^{\beta}. \tag{3.2}$$

The final term is the additional benefit from the imitation opportunity. $\beta = 0$ eliminates firm-to-firm ability transmission. At its other extreme, $\rho = \beta = 1$ simplifies the law of motion to $z' = e^{c+\varepsilon} \max\{z, \hat{z}\}$. Our goal is, in part, to let the data tell us where these two parameters fall.[12]

We assume learning occurs via operating firms. Denoting M^f as the equilibrium distribution of operating firms (defined formally below), we write $\widehat{M}(\hat{z}; M) =$

[11] The constant returns assumption for suppliers eliminates any strategic interactions between supplier-firm bargaining games.

[12] The max operator in the diffusion process rules out any benefit to the higher ability agent in the match. We show later that this is not a critical assumption but does have some empirical support (Brooks et al., 2018; Jarosch et al., 2021). Since our empirical results will eventually imply it, we maintain it throughout.

$M^f(\hat{z}; M)^{1/(1-\theta)}$.

Occupational Choice Decision and Source Distribution The aggregate state of the economy is the distribution of ability, $M(z)$. At the start of each period, each agent chooses to either operate a firm or engage in wage work given her ability z and aggregate state M.

The flow utility of operating a firm is

$$\begin{aligned} u^f(z, M) &= \max_{s,x,n \geq 0} \; \omega \log(x^\alpha n^\eta - p_x x - wn) + (1-\omega)\log(1-s) \\ s.t. \quad & m = f(s, z) \\ & p_x = \text{argmax}_{p_x} \; [\pi]^\nu [\pi^s(m)]^{1-\nu} \end{aligned}$$

The first constraint is the type of supplier met with search effort s. The second guarantees that the realized cost is the outcome of Nash bargaining between the firm and its supplier. Flow utility for a worker is $u^w(z, M) = \omega \log(w(z, M))$, where w is the equilibrium wage. We summarize the occupational choice decision as $u(z, M) = \max\{u^f(z, M), u^w(z, M)\}$. We denote decision rule $\phi(z, M) = 1$ as firm operation and $\phi(z, M) = 0$ as wage work.

The source distribution is therefore

$$\widehat{M}(\hat{z}; M) = \left(\frac{\int_0^{\hat{z}} \phi(z, M) dM(z)}{\int_0^{\infty} \phi(z, M) dM(z)} \right)^{\frac{1}{1-\theta}}.$$

Recursive Formulation Taken together, the value of entering the period with ability z and aggregate state M is

$$\begin{aligned} v(z, M) &= \max\{u^f(z, M), u^w(z, M)\} + (1-\delta) \int_\varepsilon \int_{\hat{z}} v(z'(\hat{z}, \varepsilon; z), M') \widehat{M}(d\hat{z}, M) dF(\varepsilon) \\ s.t. \quad & z'(\hat{z}, \varepsilon; z) = e^{c+\varepsilon} z^\rho \max\left\{1, \frac{\hat{z}}{z}\right\}^\beta \end{aligned}$$

Equilibrium We study the stationary equilibrium of this model, which includes the value function v, decision rules for occupation $\phi(z, M)$, effort $s(z, M)$, labor $n(z, M)$, and intermediates $x(z, M)$, bargaining outcomes $p_x(z, M)$, and a distribution of ability $M(z)$, such that the value functions solve the agent's problem above with the associated decision rules, and the aggregate state evolves according to

$$\begin{aligned} & M'(z') := \Lambda(M(z')) \\ = \quad & \delta G(z') + (1-\delta)\int_0^{\infty}\int_0^{\infty} F(\log(z') - \rho\log(z) - \beta\log(\max\{1, \hat{z}/z\}) - c)\widehat{M}(d\hat{z}; M)M(dz). \end{aligned}$$

$\Lambda(M)$ is the law of motion for the aggregate state. It consists of the ability of new entrants, $\delta G(z)$, and the evolution of ability for surviving agents. In the stationary equilibrium, $M^*(z) = \Lambda(M^*(z))$.

3.2 Characterization of Equilibrium

We summarize a few useful features of the equilibrium in the following proposition, with the proof in the Appendix.

Proposition 1. *The following results hold in the equilibrium of the model:*

1. *The solution to the Nash bargaining game is a constant markup over marginal cost, $p_x(m) \propto m$.*
2. *The equilibrium profit function can be written as $\pi(z, w) = A(w)z$, where $A(\cdot)$ depends only on the equilibrium wage and parameters of the model.*
3. *All agents with $z \geq \underline{z}$ operate firms. That cut-off is given by the function $\underline{z}(w) = w^{\frac{1-\alpha}{1-\eta-\alpha}}C$ for a constant C.*

The first two results of Proposition 1 turn out to be useful for calibration below.[13] The cut-off occupational choice rule follows directly from the result that $\pi(z, w) = A(w)z$. This occupational choice margin is central to the diffusion externality in the model. Since agents do not take into account how their occupational choice affects the learning of others, the decentralized economy allocates more agents to firm operation than the planner would (we solve the planner's problem in Appendix E). We isolate the importance of this margin in the quantitative results below.

3.3 Model Calibration

We now turn to parameterization the model. This involves two steps. In the first step we estimate the diffusion parameters. The beauty of the partial equilibrium RCT is that we can estimate these parameters independent of the remaining structure of the model. The second step is a more standard calibration given these diffusion

[13] That $\pi(z, w) = A(w)z$ follows because matching between suppliers and firms is deterministic. Uncertainty these matches is straightforward to include. It requires an adjustment to the diffusion parameter estimation analogous to a measurement error adjustment. We discuss this measurement error extension in the Appendix.

parameters. The full set of moments and parameter values are reported in Table 2 in Appendix A, and the model matches them well.

3.3.1 Estimating Diffusion Parameters

When we estimate the diffusion parameters we restrict attention to the baseline and survey wave 1 quarter post-treatment. There are two sub-steps that we describe in turn.

Step 1: Estimating (β, ρ) from Treatment Data We first estimate the parameters β and ρ. These control the change in ability conditional on a given match $\hat{z}$ from the law of motion (3.2). We can re-write this law of motion in terms of profit because $z \propto \pi$ in equilibrium,

$$\log(\pi') = \tilde{c} + \rho \log(\pi) + \beta \log\left(\max\left\{1, \frac{\hat{\pi}}{\pi}\right\}\right) + \varepsilon \tag{3.3}$$

where π and $\hat{\pi}$ represent an individual and her match's profit, and $\tilde{c}$ is a constant.

If we consider just treated firms (and denote that set $\mathbf{T}$) and note that in our RCT $\hat{\pi}_i \geq \pi_i \;\; \forall i \in \mathbf{T}$, we have

$$\log(\pi_i') = \tilde{c} + \rho \log(\pi_i) + \beta \log\left(\frac{\hat{\pi}_i}{\pi_i}\right) + \varepsilon_i \qquad \text{for } i \in \mathbf{T}. \tag{3.4}$$

Equation (3.4) is a straightforward linear regression run only on treated firms. These coefficients $(\hat{\beta}, \hat{\rho})$ are equal to their structural counterparts. This relies critically on the fact that matches within the treatment are randomized; the shocks ε need not be i.i.d., and therefore require additional assumptions so that coefficient estimates are not biased.

Those results are in Column (1) of Table 1, and we find that $\beta = 0.538$ and $\rho = 0.595$. For some sense of what these parameters represent, note that the closed form solution for $\hat{\beta}$ is

$$\hat{\beta} = \frac{cov(\log(\pi_i'), \log(\hat{\pi}_i))}{\sigma^2_{\log(\hat{\pi})}}.$$

That is, it measures the covariance between firm i's match quality as measured by her match's baseline profit $\hat{\pi}_i$ and i's profit gains π_i'. It is then normalized by the exogenous variation fed into the experiment, $\sigma^2_{\log(\hat{\pi})}$. The result $\hat{\beta} = 0.538$ implies there is a strong, but far from perfect, internalization of match knowledge.

ρ measures the persistence of profit after properly controlling for different match quality across individuals. Here, we find that profit is not very persistent, at least relative to measures in richer countries or among larger firms. This is consistent with the mechanisms discussed in the previous section.

Step 2: Estimating θ from Relative Profit Once we have (β, ρ) we can estimate θ. The relevant empirical moment to estimate θ is the average treatment effect. We measure the average treatment effect in percentage terms, so that $ATE^{data} := \mathbb{E}[\pi'_T]/\mathbb{E}[\pi'_C]$. After some algebra, we can write the model's implied average treatment effect as

$$ATE^{model}(\theta) = \frac{\int\int \pi^\rho \max\left[1, \hat{\pi}/\pi\right]^\beta d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^\rho \max\left[1, \hat{\pi}/\pi\right]^\beta dM^f_\pi(\hat{\pi})^{1/(1-\theta)} dH_{C,\pi}(\pi)} \tag{3.5}$$

where $H_{T,\pi}$, $H_{C,\pi}$ are the empirical cdfs of baseline profit for treated and control firms and $\widehat{H}_{T,\pi}$ is the empirical cdf of the matches for treated firms. $ATE^{model}(\theta)$ is the average treatment effect that would result from running the exact same RCT in the model economy that was discussed in Section 2.

Once we have estimates (β, ρ), equation (3.5) includes only one unknown: θ. The remaining terms include only distributions of profits. For a given RCT, equation (3.5) is monotone and continuous in θ. This means that a solution $ATE^{data} = ATE^{model}(\theta)$ is unique if it exists. A solution exists if $ATE^{data} \in [\Gamma^{min}, \Gamma^{max}]$, where these bounds are given by

$$\begin{aligned}
\Gamma^{min} &= \inf_\theta \frac{\int\int \pi^\rho \max\left\{1, \hat{\pi}/\pi\right\}^\beta d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^\rho \max\left\{1, \hat{\pi}/\pi\right\}^\beta dM^f_\pi(\hat{\pi})^{1/(1-\theta)} dH_{C,\pi}(\pi)} \\
\Gamma^{max} &= \sup_\theta \frac{\int\int \pi^\rho \max\left\{1, \hat{\pi}/\pi\right\}^\beta d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^\rho \max\left\{1, \hat{\pi}/\pi\right\}^\beta dM^f_\pi(\hat{\pi})^{1/(1-\theta)} dH_{C,\pi}(\pi)}.
\end{aligned}$$

These bounds guarantee the treatment effect can be rationalized by the model. For example, if $\beta = 0$ the model cannot rationalize any non-zero treatment effect. This implies $\Gamma^{min} = \Gamma^{max} = 0$. Usefully, these bounds are computable with the relevant RCT data and estimates (β, ρ). Our treatment effect falls within the bounds and therefore there is a unique θ such that $ATE^{data} = ATE^{model}(\theta)$.

The relevant empirical moment is in Column (2) of Table 1. This implies $\theta = -0.417$. For a sense of magnitude, this implies that the expected match ability is 95 percent of the ability of the average operating firm. Matches are slightly worse than random draws from the operating firm distribution.

3.3.2 Calibration of Remaining Parameters

We next calibrate the remaining parameters conditional on the already-estimated diffusion parameters. Since the diffusion parameters do not make use of any of the

Table 1: Moments for Diffusion Parameter Estimation

	(1)	(2)
β	0.538	
	(0.273)**	
ρ	0.595	
	(0.273)**	
Treatment		891.990
		(280.720)***
R^2	0.053	0.047
Control Avg	–	1897.851

Table notes: Standard errors are in parentheses. The top and bottom one percent of dependent variables are trimmed. Statistical significance at 0.10, 0.05, and 0.01 is denoted by *, **, and, ***.

remaining parameters of the model, we do not need to adjust them based on what we calibrate here. We make use of the baseline field data that is a representative sample of firms in Dandora, Kenya to calibrate to the local economy.

We assume both exogenous shock processes are lognormal, so that new entrants draw from $G \sim logN(\mu_0, \sigma_0)$ and existing firms from $F \sim logN(\mu, \sigma)$. We normalize $\mu_0 = 0$. We note that the drift in the ability law of motion c and the mean of the exogenous shocks for existing firms, μ (from c.d.f. F), are not separately identified. We set $\mu = -\sigma^2/2$ so that $\mathbb{E}[e^{\varepsilon}] = 1$. This leaves 8 remaining parameters.[14]

The remaining 8 parameters can be broken into two groups. The first group is matched one-to-one with a given moment or value (δ, σ_0, ν). The death rate δ is set average age of population in the study, which is 34.[15] The standard deviation of new entrant ability matches the variance of log profit for firms that have been open for less than 3 months, which implies $\sigma_0 = 1.00$. Finally, we note that given the model set-up the bargaining power of firms in its supplier negotiation (ν) has no effect the results. Thus, we set it as $\nu = 0.5$ for simplicity.

This leaves 5 parameters – σ, c, ω, α, and η – which we target to jointly hit 5 moments. While jointly calibrated, each has an intuitive counterpart. We choose σ to match the standard deviation of log profit in the economy, equal to 0.99. The drift parameter c in ability is targeted to match the average profit of all firms relative to those who entered less than one year ago (1.51). The utility parameter ω is set to match the share of employment in wage work. The most recent Kenyan census (via IPUMS, 2020) implies that 48 percent of employment in Embakasi Constituency (the local area in Kenya that includes our study site) is in wage work.

[14] On the utility side, they include the relative weight of consumption ω and the agent death rate δ. The remaining parameters dealing with technology and ability evolution are the parameters α and η, the ability growth term c, and the standard deviations of the exogenous shocks σ_0 and σ. The final parameter is the bargaining weight ν.

[15] The constant death rate δ implies a geometrically distributed age distribution with mean $1/\delta$ and, assuming a new agent is 18 years old, implies an average age of 64 quarters in the model. We match actual age rather than age of the firm because agents can move between firm operation and wage work during their lifetime.

α and η are then set to match two moments. The first is the standard deviation of the log unit intermediate cost across firms. After removing industry fixed effects that may bias those estimates, the standard deviation in the data is 1.61. The second moment is the wage bill relative to intermediate spending. The average firm at baseline has a ratio of 0.13. Our Cobb-Douglas assumption implies that $wn/p_x x = \eta/\alpha$, so we set $\eta = 0.13\alpha$. The standard deviation of intermediate cost is then exactly pinned down by the cost ratio and the variability in log profit. This follows from the fact that Proposition 1 implies $\sigma_{\log(p_x)} = \left(\frac{1-\alpha-\eta}{\alpha}\right)\sigma_{\log(\pi)}$. Plugging $\eta = 0.13\alpha$ yields $\alpha = \frac{\sigma_{\log(\pi)}}{\sigma_{\log(p_x)}+1.13\sigma_{\log(\pi)}}$.

The complete list of moments and parameter values are in Table 2.

Table 2: Targets and Parameter Choices

Model Parameter	Description	Parameter Value	Target Moment	Source	Target Value	Model Value
Group 1	*Diffusion Block from RCT*					
β	Intensity of diffusion	0.538	Estimated parameter from regression (5.5)	RCT results	0.538	0.538
ρ	Persistence of ability	0.595	Estimated parameter from regression (5.5)	RCT results	0.595	0.595
θ	Match technology "quality"	-0.417	Treatment effect in quarter 2 (as % above control)	RCT results	0.403	0.403
Group 2	*Matched one-to-one with parameter*					
δ	Death rate of firms	0.016	Average age of baseline business owners	Baseline survey	0.09	0.09
σ_0	St. dev. of new entrant ability distribution	0.961	Variance of log profit among new entrants	Baseline survey	1.00	1.00
ν	Firm bargaining weight	0.50	Set exogenously	–	–	–
Group 3	*Jointly targeted*					
σ	St. dev. of exogenous ability shock distribution	0.75	Standard deviation of log profit in all firms	Baseline survey	0.99	0.99
c	Growth factor in ability evolution	-1.92	Ratio of average profit of all firms to new entrants	Baseline survey	1.51	1.51
ω	Consumption utility weight	0.53	Fraction of employment in wage work	IPUMS	0.48	0.48
α	Ability elasticity in supplier search	0.36	Standard deviation of log inventory cost	Baseline survey	1.61	1.61
η	Ability elasticity in supplier search	0.05	Average cost ratio	Baseline survey	0.13	0.13

Table notes: Group 1 is jointly chosen from the experimental data. Group 2 are also set to match baseline data moments, but match 1-1 with target moments. Parameters in Group 3 are calibrated to jointly match moments.

3.4 Dynamic Treatment Patterns

Our calibration does not use the full time series of the average treatment effect, only the data before and the first period after treatment. Before turning to the quantitative results we ask if the model can replicate the observed pattern. Figure 2a replicates the time series of the ATE over 5 quarters in the model. The first two quarters are matched by construction. Both the model and data predict no treatment effect by $t = 3$. The model under-predicts the effect in $t = 2$ so, if anything, the model understates the partial equilibrium RCT dynamics.

Figures 2b and 2c show how the diffusion parameter estimates β and ρ inform our model prediction. If we instead estimated a high knowledge persistence (high ρ) and low ability to learn (low β), we could have instead predicted a 20 percent treatment effect at $t = 5$. Thus, the parameter estimates from the RCT put discipline on the fade-out pattern we predict from the model.

Figure 2: Relationship between Treatment Persistence and Diffusion Intensity β

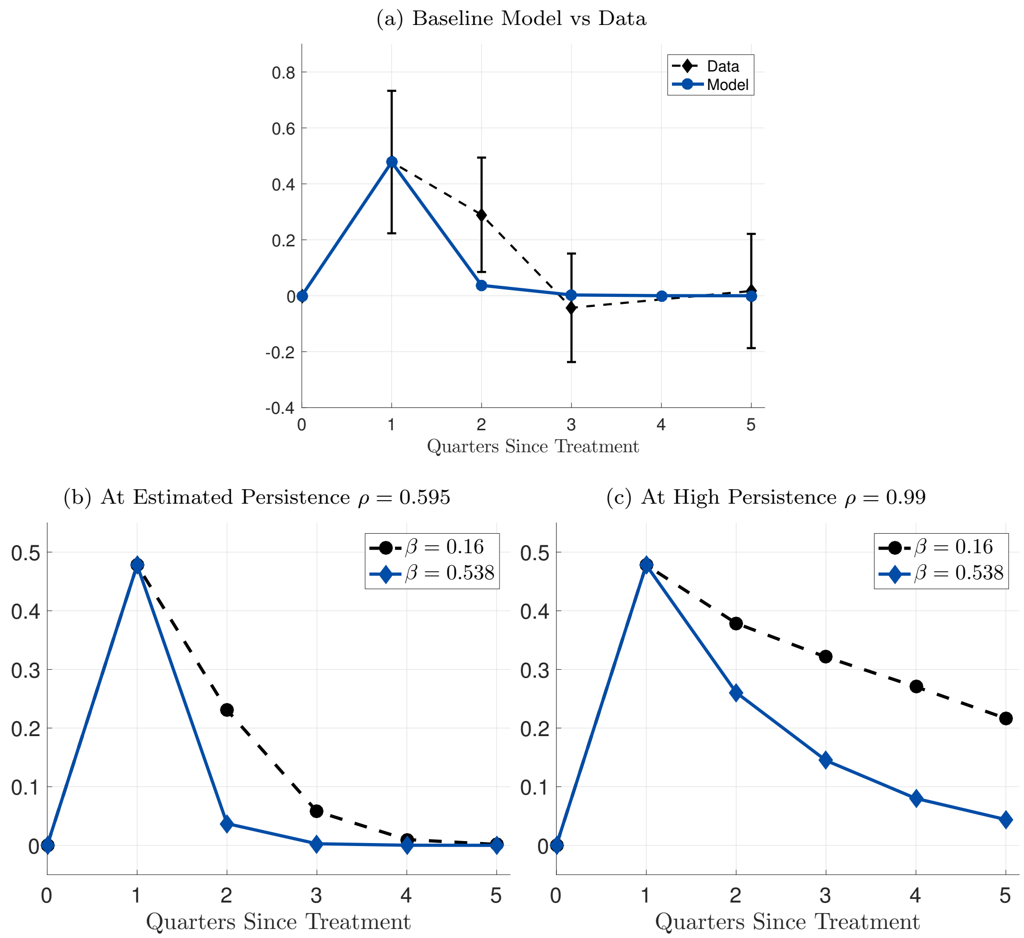

(a) Baseline Model vs Data

(b) At Estimated Persistence $\rho = 0.595$

(c) At High Persistence $\rho = 0.99$

As we will discuss below, the parameter β is critical for generating general equilibrium gains from diffusion. Yet higher β also hastens fade-out. An immediate consequence of the results here is that treatment effect persistence need not be positively correlated with the equilibrium gains that could be achieved at scale. This difficulty in linking average treatment effect moments to at-scale equilibrium implications is a theme we will revisit in the quantitative results as well.

4 Quantitative Results

With the calibrated model in hand, we now turn to quantitative results. The goal is to understand the equilibrium implications of better matching technologies. We do so by varying θ. The intervention discussed in Section 2, along with those in the introduction, suggest that at least part of the meeting frictions we impose via θ are policy-malleable. We think of this exercise as measuring the gains available to a policymaker who adjusts this parameter by making it easier to find high-quality matches.

4.1 Impact of Better Matches in Equilibrium

Quantitative Experiment We implement our "better matching equilibrium" by shocking θ so that it moves 25 percent closer to its limit of $\theta = 1$. This increases baseline $\theta = -0.417$ to $\theta^{new} = -0.063$.[16] We study the long run, general equilibrium change in income that results by comparing the two stationary equilibria.

Quantitative Outcomes Aggregate moments are reported in Table 3. The first column is the baseline economy ($\theta = -0.417$) with aggregate moments from the stationary equilibrium normalized to one. The second column presents a new equilibrium with $\theta^{new} = -0.063$ but fixes the wage at the baseline level. Column three allows the wage to adjust as well.

Overall, income rises by 11 percent. This is made up of two general equilibrium effects. The first is that the new matching technology directly affects the ability distribution by making it easier to learn from high ability agents. The second is an amplification effect through prices.

[16] The choice of 25 percent is somewhat arbitrary, but also irrelevant for our results net of some differences in magnitudes. In practice, policy changes are often motivated by experiments but are not exact replications for administrative reasons. We think of this as a policymaker creating a program that has the same spirit as the smaller scale intervention. This could be through extension programs, better use of information technology, or explicit mentoring programs, among others. Regardless, in the Appendix we provide the results when the aggregate policy exactly replicates the estimated RCT gains. The same results hold with larger magnitudes.

Table 3: Equilibrium Moments

	(1) Baseline	(2) Fixed Wage	(3) New Equilibrium
Income	1.00	1.07	1.11
Ability	1.00	1.08	1.12
Aggregate Wage-Labor Supply	1.00	0.92	0.98
Wage	1.00	1.00	1.13

Table notes: Column (1) is the initial equilibrium, normalized to one. (2) and (3) report the new stationary equilibrium after shocking the matching technology, where (2) holds the wage fixed at its baseline level and (3) allows it to adjust.

We decompose the relative importance of these two channels in columns 2 and 3. Column 2 isolates the direct effect on ability. Average ability rises by 7 percent and the labor supply declines by 8 percent as the distribution shifts mass across the (fixed) cut-off ability level that defines occupational choice (see Proposition 1). Average income rises by 7 percent.

Column 3 allows the equilibrium wage to adjust. The wage increases by 13 percent as the ability to find lower-cost suppliers increases the marginal product of labor. This competitive pressure causes the lowest z firms to exit and work for a wage instead. Removing relatively low quality firms allows for easier learning. This amplifies the direct effect on ability through the diffusion process.

Of the total increase in ability, 67 is from the direct effect and 37 is from the amplification through prices. A similar magnitude holds for income: 64 percent is from the direct effect, while 36 is through the price adjustment.

4.2 How much discipline does the ATE provide?

The average treatment effect is often used to determine the success of an intervention. A separate, but related, issue is whether or not it disciplines the gains that could be achieved at scale. In this section, we show how the various moments from the RCT jointly provide discipline to the quantitative results, and how their relative importance differs. The main result here is that heterogeneity *across* treated agents matters more than the average impact.

To investigate the importance of the average treatment effect, we fix it at the realized empirical level ATE^{data} and continue to use it to discipline the model. However, instead of using the within-treatment heterogeneity to estimate β and ρ, we simply assume values for them and ask how the results change. Thus, the model will remain consistent with our observed ATE but not with the covariance within the treatment.

Practically, this requires re-estimating θ given the (counterfactual) values of β and ρ. We then re-compute the gains by increasing θ by 25 percent from each baseline value.

Results Figure 3a shows the implied θ for each assumed (β, ρ) combination. Counterfactually lower β and ρ require a lower θ to match ATE^{data}. Intuitively, what this means is that firms now gain little from a good match. To match the observed average treatment effect, the model requires that the treatment is more valuable to firms. It does so by inferring control firms receive worse baseline matches. This increases the difference in average match quality between the fixed treatment and control firms, guaranteeing the the treated firms gain more. Put more succinctly, the model trades off intensive margin gains for extensive margin gains.

Figure 3b shows the aggregate gains from a 25 percent increase in matching quality. Average income increases anywhere between 0.6 to 38 percent, a 63-fold range, by varying only the importance of (β, ρ) relative to θ for the same ATE. There are two main takeaways here. First, even conditional our observed ATE, the model could have predicted either very small or very large equilibrium gains. Thus, there is nothing specific about our model that guarantees one or the other; it depends on which way the empirics push the parameters. Second, a large amount of the discipline on the equilibrium gains comes from moments other than the average treatment effect. It depends critically on additional moments derived from how the intervention's impact varies across individuals in the treatment group.

A different way to see this result is to ask which average treatment effects are consistent with an 11 percent increase in aggregate income. To do so, we replicate the same exercise above for every average treatment effect between 0.01 and 70 percent. That is, we assume parameters (β, ρ) then choose θ to match the relevant ATE. The results are in Figure 4. The previous results, in Figure 3 are a vertical slice of this figure. That is, the observed ATE pins down a vertical slice of this band of outcomes, where the width of the band is given by how the results vary for (β, ρ) combinations. The main takeaway from Figure 4 is that our 11 increase in income is consistent with nearly any average treatment effect.

Overall, the results show that caution is required when attempting to extract policy recommendations from empirical moments here. We offer a way to estimate different empirical moments that help guide the mapping from observed treatment effects to the gains that can be achieved at scale. We next turn to understanding why certain moments matter more than others.

Figure 3: Aggregate Implications of Varying (β, ρ) at Baseline ATE

(a) Implied θ (b) Aggregate Gains

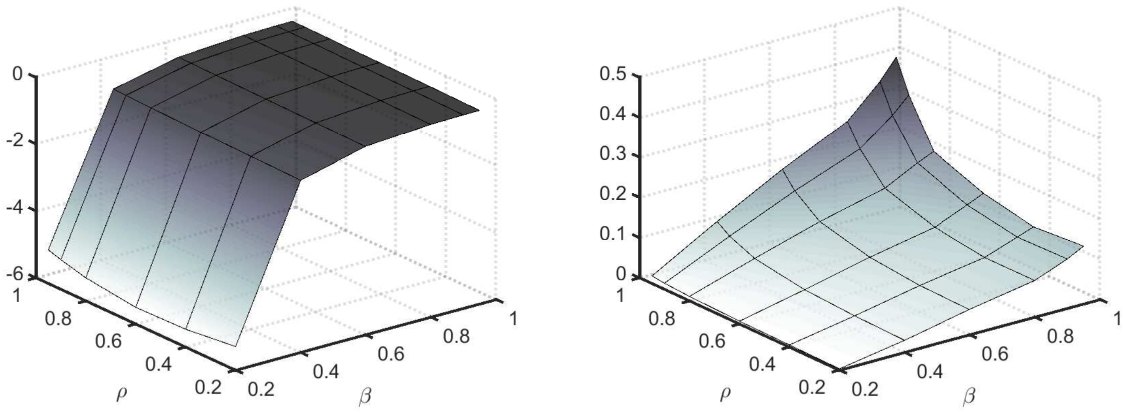

Figure notes: For each value of (β, ρ), (a) plots the implied θ that guarantees the model hits our baseline ATE. (b) then plots the implied aggregate gains from scaling the intervention as a percentage change in average income from the baseline equilibrium (i.e., $0.4 = 40\%$).

Figure 4: Relationship between average treatment effect and equilibrium increase in income

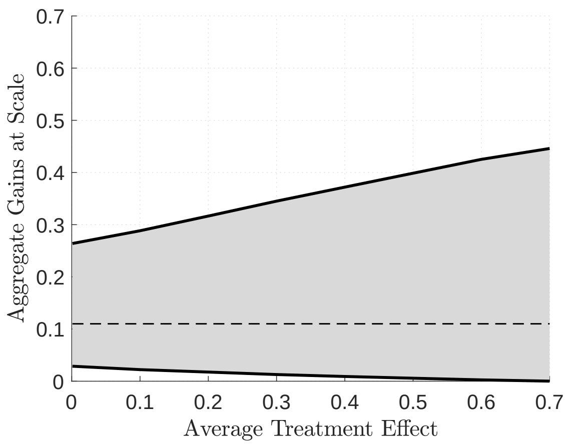

Figure notes: The shaded region shows all possible realizations of aggregate gains that can be achieved while holding the average treatment effect fixed by re-estimating the extensive margin parameter θ for each (β, ρ) combination. Multiply by 100 for percentage gains.

4.3 Understanding the Quantitative Importance of Covariance

We next turn to understanding these results, in particular, why β – governed by covariance – plays such a large quantitative role relative to θ – governed by the average effect. Our main conclusion is that the micro and macro returns are driven by different parameters.

We start by varying θ and β, holding the remaining parametrization fixed.[17] We

[17]We omit ρ in the interest of clarity, as β plays a larger role. A similar pattern emerges for ρ, though with smaller

the measure the relevant micro and macro moment. On the micro side, we measure the average treatment effect from re-running our RCT in the model at these counterfactual parameters. On the macro side, we conduct the same quantitative experiment as above. The results are in Figure 5.

Figure 5: Variation at RCT and Aggregate Levels by Parameters

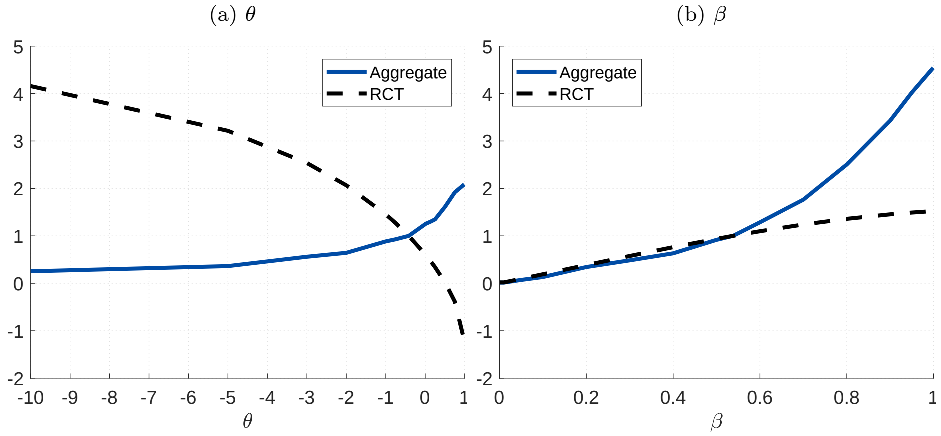

Figure notes: Each figure varies the listed parameter while holding the remaining parameterization fixed at its baseline value. "RCT" is the implied average treatment effect, while "Aggregate" is the gains from scaling the intervention as defined in the text. Results are presented as growth rates $g := (\mathbb{E}[y^{post}]/\mathbb{E}[y^{pre}]) - 1$, then normalized relative to our baseline parameterization as $g/g^{baseline}$ so that all lines cross through 1.

Two patterns emerge that show why the ATE offers little discipline on the quantitative results. First, the model-implied ATE responds strongly to changes in θ. This means that over reasonable empirical levels of the average treatment effect, there is limited scope for θ to vary. For example, Figure 5a shows that if $ATE^{data} \in (0, 100\%)$, then θ needs only to vary within $\theta \in (-0.5, 0.5)$ to match it. Therefore, if the ATE is to play a major role, it must be that the equilibrium gains vary substantially over a small range of θ. Figure 5 shows this is not the case; the equilibrium gains respond much more strongly to variation in β.

Put together, we get back to our previous result: the equilibrium gains vary widely for a given ATE because the micro-level and equilibrium gains are driven quantitatively by different parameters in the model.

The rationale is as follows. The equilibrium amplification of learning is closely connected to the mass of agents in the right tail, as heavier-tailed distributions facilitate more learning. β governs this here. To see this, Figure 6 plots how the ability distribution responds to the improved source distribution under changes for β and,

magnitudes.

for comparison, θ. Figures (a) and (b) plot the survival function $1 - M(z)$, or the fraction of the population whose ability is larger than z. Both axes are measured in logs. We do so under our baseline parameterization, then also vary β and θ to make them 50 percent closer to their theoretical maximum of one.[18] Figure 6c shows how the difference varies for each of these two parameters.

Figure 6: Equilibrium Survival Function as Parameters Change

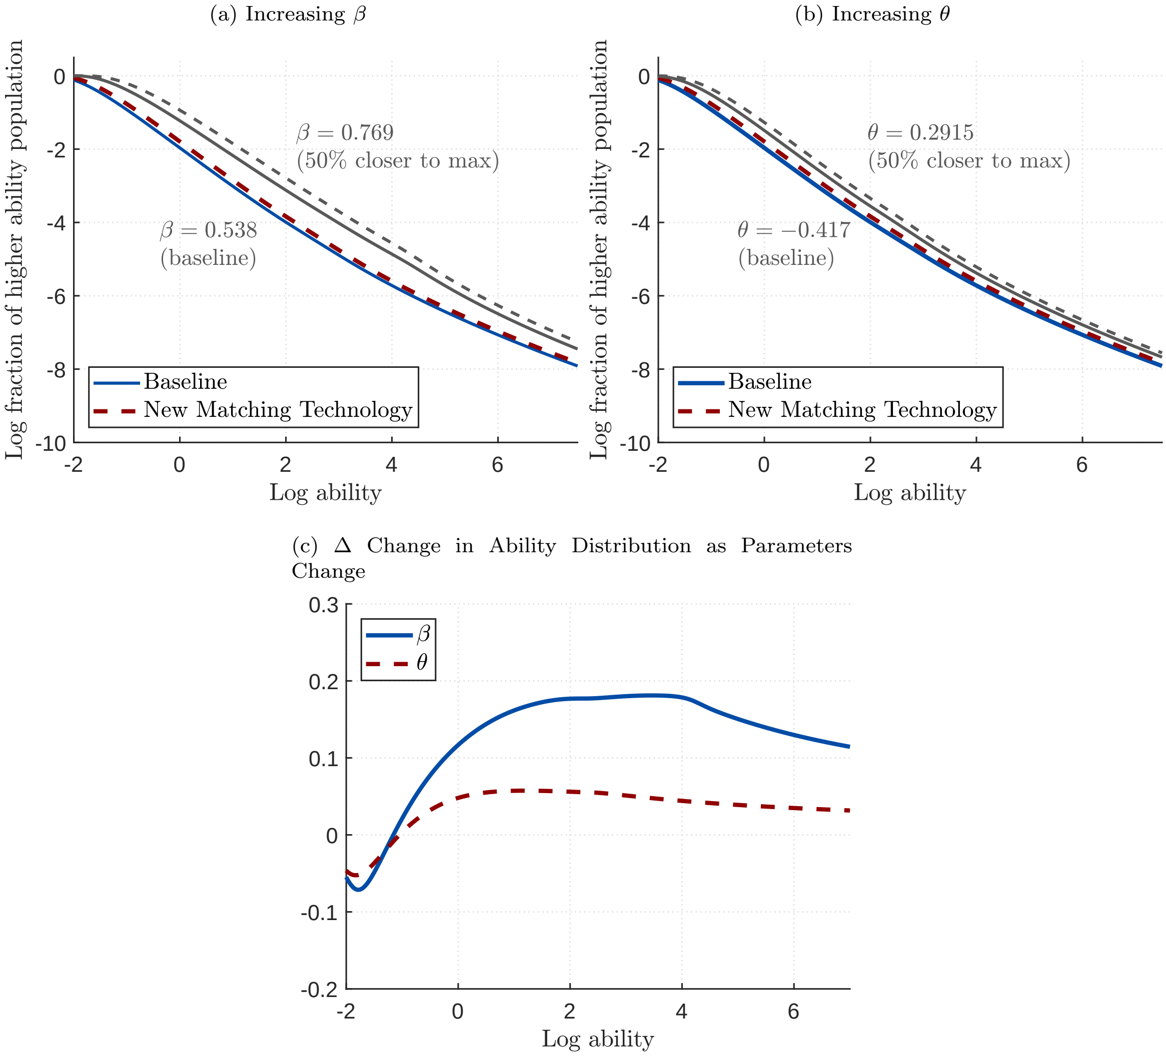

Figure notes: Figures (a) and (b) plot the ability distribution at the baseline and improved matching technology, measured as the log survival function. Figure (c) plots the difference in the mass across the ability distribution. This is measured as the difference-in-difference of the curves plotted in (a) and (b). That is, denoting $S = \log(1 - M(z))$ as the log survival function, it plots $\Delta S = (S^{new}_{x'} - S^{base}_{x'}) - (S^{new}_{x} - S^{base}_{x})$ where $x \in \{\theta, \beta\}$ and x' is the counterfactually higher level of that parameter.

[18]That is, we increase β to 0.769 because (1-0.769)/(1-0.538) = 0.5. We similarly increase θ.

This ability of the learning technology to push agents into the tail is what covariance measures in the experiment. It takes a fixed set of matches and measures how much larger the impact is for a relatively better match. The overarching conclusion here is that estimating aggregate effects requires understanding and estimating the parameters that align with the aggregate economic forces in the model. Relatedly, these results show that the aggregate importance of this policy change is not predetermined by the model. We could have predicted smaller effects with simple changes in parameter values. The discipline provided by the empirical moments plays a critical role and highlights the complementarity between the aggregate model and well-identified causal effects.

5 Broader Lessons for Understanding RCT Gains at Scale

While we focus on one specific RCT to fix ideas, Brooks et al. (2018) is part of a broader set of RCTs that create random learning opportunities for firm owners. Moreover, the promising results of these RCTs have been used as evidence for large-scale policy decisions (e.g., World Bank, 2020). In this final section, we take a step back to probe whether our results provide any more general guidance on measurement in these types of partial equilibrium RCTs, or whether they are driven by the exact model we write down or the exact intervention we use.

We ask specifically whether we can use the same moments to pin down the same diffusion parameters in a broader class of models and interventions. To do so, in this section we lay out sufficient conditions that guarantee the same procedure holds. These conditions turn out to cover a wide range of potential models and interventions. We begin by laying out the economic environment, then describe a class of interventions as a data-generating process from the model.

5.1 Class of Models

Time is discrete, and there is a population of agents with ability z. We make three assumptions on the economic environment. The first is how ability evolves.

Assumption 1. *Given ability z this period, an imitation opportunity $\hat{z}$, and a random shock ε, ability next period z' is given by*

$$z'(z, \varepsilon, \hat{z}) = e^{c+\varepsilon} z^{\rho} \max \left\{ 1, \frac{\hat{z}}{z} \right\}^{\beta}, \tag{5.1}$$

where the parameter c is a constant growth term, β is diffusion intensity, and ρ is persistence. ε is uncorrelated with z and $\hat{z}$ but not i.i.d. across agents.

This assumption mirrors our technological assumption used in the model above. The next two assumptions are on equilibrium outcomes of the model. The first relates unobservable ability z to observable characteristics.

Assumption 2. *In any period, equilibrium profit π is proportional to ability. For any two agents i and j, $\pi_i/\pi_j = z_i/z_j$.*

This result holds in our model in equilibrium. The critical feature in Assumption 2 is that there is some way to move between unobserved ability z and observable characteristics. Thus, it is straightforward to allow more general relationships that allow, for example, production function estimation. We take up extensions in the Appendix.[19]

Finally, the third assumption puts characteristics on the source distribution from which $\hat{z}$ is drawn in equilibrium. We denote the cumulative density function of $\hat{z}$ as $\widehat{M}(\hat{z}; z, \theta)$. $\widehat{M}$ implies that different z agents can draw from different distributions and those distributions depend on a technological parameter θ. The role of θ is summarized in Assumption 3.

Assumption 3. *The source distribution of imitation draws $\hat{z}$ can be characterized in equilibrium by a cumulative density function of the form $\widehat{M}(\hat{z}; z, \theta)$ with the following properties:*

1. *For every z and $\hat{z}$, $\widehat{M}$ is continuous in θ*
2. *$\theta_1 < \theta_2 \implies \widehat{M}(\hat{z}; z, \theta_2)$ first order stochastically dominates $\widehat{M}(\hat{z}; z, \theta_1)$.*

Our assumption, that $\widehat{M}(\hat{z}; z, \theta) = M^f(\hat{z})^{\frac{1}{1-\theta}}$, satisfies the continuity and FOSD requirements. But Assumption 3 is broader than our specific assumption, spans search models through pure assignment models, where θ measures the complementarity between abilities. We discuss various alternatives that fall under this assumption in the Appendix.

Together, these 3 assumptions characterize diffusion in a variety of models, including those summarized in Alvarez et al. (2008).[20] It also includes the model developed in Section 3.

[19] In addition, all of the results in this section are robust to the inclusion of unobserved idiosyncratic variation via measurement error or distortions. These extensions build off an active literature on non-linear error-in-variables models (see Schennach, 2020, for a thorough review).

[20] For example, Lucas (2009) follows by setting $\beta = \rho = 1$, $c = 0$, and making F degenerate (a special case of Assumption 1), setting $\pi = z$ (Assumption 2) and assuming $\widehat{M}(\hat{z}; z, \theta) = M^\theta$ where M is the equilibrium c.d.f. of ability and θ indexes the number of draws a firm receives each period (Assumption 3). Buera and Oberfield (2020)'s delineation between random, original innovations and learning from others is a particular interpretation of θ in Assumption 3. More examples, including those that move away from random meetings, are included in the Appendix.

5.2 Class of Interventions that Facilitate Better Matches

We now define a data-generating process that covers a class of interventions including the Brooks et al. (2018) RCT discussed in Section 2, but also includes Atkin et al. (2017), Cai and Szeidl (2018), Lafortune et al. (2018), and Fafchamps and Quinn (2018), among others.[21,22]

An important consideration as we generalize away from the specifics of Brooks et al. (2018) is the difference between an imitation opportunity and an individual in the model. In our work above, they are the same. That is, if person i meets with person j, i's imitation opportunity is $\hat{z}_i = z_j$. But these are separate notions. If, for example, i meets with a group of agents $j_1, \ldots, j_N$ and receives the average knowledge of those N firms, i's imitation opportunity is $\hat{z}_i = (1/N)\sum_{k=1}^{N} z_{j_k}$.

With that in mind, we formally define the data generating process in Assumption 4.

Assumption 4. *A set of agents with profit $H_\pi(\pi)$ are observed in two consecutive periods. That set is partitioned into two subsets $\mathbf{C}$ and $\mathbf{T}$ (i.e., "control" and "treatment"), characterized by their profit distributions $H_{C,\pi}(\pi)$ and $H_{T,\pi}(\pi)$. The following conditions hold:*

(a) The average unobserved shock for agents in $\mathbf{T}$ and $\mathbf{C}$ are equal.

(b) We cannot observe the imitation opportunity agents in $\mathbf{C}$. They are drawn from a distribution $\widehat{M}_\pi(\hat{\pi}; \pi, \theta)$.

(c) We observe the imitation opportunity for agents in $\mathbf{T}$. The distribution of those profits are denoted $\widehat{H}_{T,\pi}(\hat{\pi}) \neq \widehat{M}(\hat{z}; z, \theta)$. Furthermore, a positive measure of agents interacts with a more profitable match.

$$\int_\pi \int_{\hat{\pi}} \mathbb{1}[\hat{\pi} > \pi] d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi) > 0.$$

(d) For any arbitrary partition of $\mathbf{T}$, characterized by $H^1_{T,\pi}(z)$ and $H^2_{T,\pi}(z)$, agents in both groups draw their ε shocks from the same distribution

Assumption (a) is the usual exclusion restriction. It puts restrictions on how the unobserved characteristics vary between control and treatment. (b) formalizes the

[21] Experimental variation is not critical. Any instrument or instruments that satisfy Assumption 4 will apply similarly. Most available evidence comes from randomized controlled trials here, so we generally focus discussion in those terms.

[22] Note that Assumption 4 is not designed to be an idealized experiment to estimate a diffusion model. One could easily develop a more useful experiment than we assume here. This is not our goal. Our goal is develop a link from existing work to the models discussed in Section 5. Assumption 4 is the formalization of the way in which we tie our hands to a particular type of variation.

intuitive notion that we cannot observe individual-level control matches, but that matches continue to happen in the background of the economy. That is, the control group does not stop meeting other agents because of the intervention.

Assumptions (c) and (d) lay out what we require from the treatment. The treatment shocks the source distribution from $\widehat{M}$ to $\widehat{H}_T$. The last assumption states a second exclusion restriction within the treatment group, guaranteeing that a comparison between treated firms is similarly unbiased.

5.3 The Need for Additional Moments and Identification

Our goal is to estimate three key parameters, β, ρ, and θ. As one might suspect, it would be impossible to estimate all three a single average treatment effect moment. We show first that it is even more dire than that, because there almost always exists a way to rationalize a given ATE. To see this, define the ATE as the percentage change in profit between firms in treatment and control,

$$ATE^{data} = \frac{\mathbb{E}[\pi_i'|i \in \mathbf{T}]}{\mathbb{E}[\pi_i'|i \in \mathbf{C}]} \tag{5.2}$$

and the model's counterpart as

$$ATE^{model} = \frac{\mathbb{E}_T[\pi']}{\mathbb{E}_C[\pi']} = \frac{\int\int z^\rho \max\{1, \hat{z}/z\}^\beta \, d\widehat{H}_T(\hat{z}) dH_T(z)}{\int\int z^\rho \max\{1, \hat{z}/z\}^\beta \, d\widehat{M}(\hat{z}; z, \theta) dH_C(z)}. \tag{5.3}$$

We summarize the relationship between the two in Proposition 2 with discussion following. All proofs are in the Appendix.

Proposition 2. *For a given pair (β, ρ) there is at most one θ such that $ATE^{model} = ATE^{data}$. If $ATE^{data} \in [\Gamma^{min}, \Gamma^{max}]$ that θ exists, is unique, and solves*

$$ATE^{data} = \frac{\int\int \pi^\rho \max\{1, \hat{\pi}/\pi\}^\beta \, d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^\rho \max\{1, \hat{\pi}/\pi\}^\beta \, d\widehat{M}_\pi(\hat{\pi}; \pi, \theta) dH_{C,\pi}(\pi)}. \tag{5.4}$$

If $ATE^{data} \notin [\Gamma^{min}, \Gamma^{max}]$ then no such θ exists. The bounds are given by

$$\begin{aligned} \Gamma^{min} &= \inf_\theta \frac{\int\int \pi^\rho \max\{1, \hat{\pi}/\pi\}^\beta \, d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^\rho \max\{1, \hat{\pi}/\pi\}^\beta \, d\widehat{M}_\pi(\hat{\pi}; \pi, \theta) dH_{C,\pi}(\pi)} \\ \Gamma^{max} &= \sup_\theta \frac{\int\int \pi^\rho \max\{1, \hat{\pi}/\pi\}^\beta \, d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^\rho \max\{1, \hat{\pi}/\pi\}^\beta \, d\widehat{M}_\pi(\hat{\pi}; \pi, \theta) dH_{C,\pi}(\pi)}. \end{aligned}$$

Proposition 2 is a generalization of our estimation procedure in Section 3.3.2. It formalizes that there are a continuum of economies that generate identical ATEs, and differ only in the importance of β and ρ relative to θ. While the proof is in the Appendix, the intuition follows our discussion in the previous section. The extensive margin of matching (governed by θ) can always be set to counteract the direct effect of a match (governed by β and ρ).[23] Bias in one dimension therefore creates bias in the other, though remains consistent with the observed ATE. Yet we know from Section 4 that aggregate outcomes are not independent of their relative importance.

However, if one can identify (β, ρ), then there is a unique θ corresponding to ATE^{data}. The key insight here is that fixing exogenously the set of treatment draws allows us to infer what must be happening to control firms, *if* we already know how a given match affects knowledge. This is the part of the technology governed by β and ρ. Identifying these parameters can be done with variation within treatment matches.

Proposition 3. *The parameters (β, ρ) are identified by coefficients from the regression run only on treated agents (all $i \in \mathbf{T}$)*

$$\log(\pi_i') = \widetilde{c} + \rho \log(\pi_i) + \beta \log\left(\max\left\{1, \frac{\hat{\pi}_i}{\pi_i}\right\}\right) + \varepsilon_i, \tag{5.5}$$

where $\widetilde{c}$ is the constant equal to the technological parameter c if $\pi = z$.

Like before, if $\hat{\pi}_i > \pi_i$ for all $i \in \mathbf{T}$, then

$$\widehat{\beta} = \frac{cov(\log(\pi'), \log(\hat{\pi}))}{\sigma^2_{\log(\hat{\pi})}}.$$

That is, $\widehat{\beta}$ measures the covariance between a firm's *ex post* profit and its match's *ex ante* profit, normalized by the match-level variation fed into the experiment exogenously. That same intuition holds continues to hold if $\hat{\pi}_i < \pi_i$ for some i, but introduces the covariance matrix into $\widehat{\beta}$ to account for the collinearity of π_i and $\max\{1, \hat{\pi}_i/\pi_i\}$.[24] Thus, this idea that we can measure additional moments that matter for at-scale gains in the empirics holds more broadly.

Propositions 2 and 3 provide a way to move between the intervention-level data and diffusion parameters in a broad class of models. It requires two straightforward linear regressions that can be run on the data collected from a class of partial equilibrium experiments. Our interpretation of these results is the following. There is substantial evidence that diffusion plays an important role in the development process, derived from both micro and macro sources. Changing learning opportunities at

[23] This is true up to the bounds Γ^{min} and Γ^{max}. These bounds exist to guarantee that the treatment effect stays within the set of values that the model can possibly rationalize.

[24] Assumption 4(c) – that a positive measure of agents benefit from the treatment – guarantees the independent variable matrix in regression (5.5) has full rank, while Assumption 4(d) guarantees the parameter estimates are not biased by unobservable shocks ε.

scale requires understanding how to use these RCT estimates to build evidence. We show that regardless of the specifics of the exact economy one has in mind, there exist simple additional moments that can be measured in these RCTs that help discipline what we should expect to find at scale.

5.4 Further Discussion and Additional Appendix Results

Before concluding, we discuss some additional theoretical and quantitative results that are available in the Appendix.

In terms of the generalized procedure developed in Section 5 to determine diffusion parameters, we show that there are additional features that can be added. First, we allow a more general law of motion for ability which can be estimated semi-parametrically with the same data generating process. Second, and relatedly, the results can be extended to include heterogeneous returns by characteristics. Third, we extend the results to a more general relationship between observables and ability that allows for the integration of production function estimation. Fourth, we show the procedure can be adjusted to include measurement error or idiosyncratic distortions that affect the relationship between profit and ability, building off an active literature on non-linear error-in-variables models (see Schennach, 2020, for a thorough review). These various adjustments introduce important considerations and complications, such as issues of power for non-parametric estimation or the introduction of the required deconvolution methods to deal with measurement error. The underlying economic intuition holds without much change however, and we focus in the main text on this relatively parsimonious set of parameters for this reason.

In terms of the quantitative results in Sections 3 and 4, we also add a number of extensions in the Appendix. First, we quantify the bias induced by mis-measured profit. Our results are a lower bound on the equilibrium gains. Idiosyncratic distortions or non-deterministic supplier search naturally generate a similar, but not identical, result. Second, we solve the social planner's problem to study the gains that could be achieved by transitioning the economy to its efficient allocation. We find that welfare gains are largest at intermediate levels of β. The results deliver a similar message as the main text: measuring the various forces that inform aggregate gains is critical. Third and finally, we use the work of Cai and Szeidl (2018) to apply our results to a different experiment. This experiment offers a useful comparison to our results in the main text: they find that gains from matches persist longer than in our baseline RCT, and also provide direct evidence of diffusion within matches. We replicate the persistence of their treatment effect. The reason is because our estimation procedure infers different parameters from different patterns of covariance within treated firms.

6 Conclusion

We develop a model to study the cost of frictions that limit potentially profitable interactions between firms. We discipline the model by linking it to a promising set of micro-level interventions that offer these same opportunities at a smaller scale. Using evidence from an RCT in Kenya, we find that equilibrium income rises by 11 percent off an average treatment effect of 19 percent. The discipline on that 11 percent increase in income comes primarily from moments other than the average treatment effect. Our results point to other moments that do. In particular, we show that a covariance moment plays a critical role in generating equilibrium gains at scale. This moment can be estimated with partial equilibrium RCT results, thus providing policy-relevant information on scalability. The results highlight the important complementarity between causal interventions and aggregate models (Buera et al., 2022).

That is not to say this single additional moment captures *everything* that matters for aggregate outcomes. This depends on the exact specifics of the economy. We leave out interesting features that may arise, such as the possibility that firms become unwilling to share information due to competition, though some simple version of such theories are feasible under our framework (see the Appendix for an example of bargaining over match surplus). One could of course always write down a more complicated model with such a feature. That part is easy. But more parameters demand more empirical moments if the goal is to estimate this process with a methodology similar to our diffusion parameter estimation. This requires more subtly and targeted variation in designing interventions. Different field experiments, designed with an eye toward aggregate theory, could further refine our understanding of key aggregates governed by a number of difficult-to-measure elasticities that affect diffusion at scale.

References

Alvarez, Fernando E., Francisco J. Buera, and Robert E. Lucas, "Models of Idea Flows," June 2008. NBER Working Paper 14135.

Atkin, David, Amit K. Khandelwal, and Adam Osman, "Exporting and Firm Performance: Evidence from a Randomized Experiment," *Quarterly Journal of Economics*, 2017, *132* (2), 551–615.

Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and A. Mushfiq Mobarak, "Can Network Theory-based Targeting Increase Technology Adoption?," *American Economic Review*, 2021, *111* (6), 1918–1943.

Benhabib, Jess, Jesse Perla, and Christopher Tonetti, "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," *Econometrica*, 2021, *89* (5), 2261–2301.

Berguist, Lauren, Benjamin Faber, Matthias Hoelzlein, Edward Miguel, and Andres Rodriguez-Claire, "Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda," April 2019. Working Paper.

Bianchi, Nicola and Michela Giorcelli, "The Dynamics and Spillovers of Management Interventions: Evidence from the Training within Industry Program," *Journal of Political Economy*, 2022, *130* (6), 1630–1675.

Blattman, Christopher and Laura Ralston, "Generating employment in poor and fragile states: Evidence from labor market and entrepreneurship programs," June 2017. Working Paper.

Breza, Emily, Arun Chandrasekhar, Benjamin Golub, and Aneesha Parvathaneni, "Networks in economic development," *Oxford Review of Economic Policy*, 2019, *35* (4), 678–721.

Brooks, Wyatt, Kevin Donovan, and Terence R. Johnson, "Mentors or Teachers? Microenterprise Training in Kenya," *American Economic Journal: Applied Economics*, 2018, *10* (4), 196–221.

Buera, Francisco J. and Ezra Oberfield, "The Global Diffusion of Ideas," *Econometrica*, 2020, *88* (1), 83–114.

_ **and Robert E. Lucas**, "Idea Flows and Economic Growth," *Annual Review of Economics*, 2018, *10*, 315–345.

_ **, Joseph P. Kaboski, and Robert M. Townsend**, "From Micro to Macro Development," *Journal of Economic Literature, forthcoming*, 2022.

_ **,** _ **, and Yongseok Shin**, "The Macroeconomics of Microfinance," *Review of Economic Studies*, 2021, *88* (1), 126–161.

Cai, Jing and Adam Szeidl, "Interfirm Relationships and Business Performance," *Quarterly Journal of Economics*, 2018, *133* (3), 1229–1282.

Caliendo, Lorenzo, Ferdinando Monte, and Esteban Rossi-Hansberg, "The Anatomy of French Production Hierarchies," *Journal of Political Economy*, 2015, *123* (4), 809–852.

Fafchamps, Marcel and Simon Quinn, "Networks and Manufacturing Firms in Africa: Results from a Randomized Field Experiment," *World Bank Economic Review*, 2018, *32* (3), 656–675.

Fried, Stephie and David Lagakos, "Electricity and Firm Productivity: A General Equilibrium Approach," 2021. Working Paper.

Fujimoto, Junichi, David Lagakos, and Mitch Vanvuren, "The Aggregate Effects of "Free" Secondary Schooling in the Developing World," 2021. Working Paper.

Giorcelli, Michela, "The Long-Term Effects of Management and Technology Transfers," *American Economic Review*, 2019, *109* (1), 121–152.

Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon Phillips, "Knowledge Diffusion in the Workplace," July 2018. Working Paper.

IPUMS, "Minnesota Population Center. Integrated Public Use Microdata Series, International: Version 7.3 [dataset]," Minneapolsi, MN. https://doi.org/10.18128/D020.V7.2 2020.

Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg, "Learning from Coworkers," *Econometrica*, 2021, *89* (2), 647–676.

Jones, Charles I., "R & D-Based Models of Economic Growth," *Journal of Political Economy*, 1995, *103* (4), 759–784.

Jovanovic, Boyan and Rafael Rob, "The Growth and Diffusion of Knowledge," *Review of Economic Studies*, 1989, *56* (4), 569–582.

Kaboski, Joseph P., Molly Lipscomb, Virgiliu Midrigan, and Carolyn Pelnik, "How Important are Investment Individisibilities for Development? Experimental Evidence from Uganda," 2022. Working Paper.

Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, *65* (6), 1389–1419.

Lafortune, Jeanne, Julio Riutort, and Jose Téssa, "Role Models or Individual Consulting: The Impact of Personalizing Micro-entrepreneurship Training," *American Economic Journal: Applied Economics*, 2018, *10* (4), 222–245.

Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E. Waugh, "The Welfare Effects of Encouraging Rural-Urban Migration," January 2021. Working Paper.

_ **and Martin Shu**, "The Role of Micro Data in Understanding Structural Transformation," *Oxford Development Studies, forthcoming*, 2022.

Lashkari, Danial, "Innovation Policy in a Theory of Knowledge Diffusion and Selection," 2020. Working Paper.

Lucas, Robert E., "Ideas and Growth," *Economica*, 2009, *76* (301), 1–19.

_ **and Benjamin Moll**, "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, 2014, *122* (1), 1–51.

Munshi, Kaivan, "Information Networks in Dynamic Agrarian Economies," in T. Paul Schultz and John Strauss, eds., *Handbook of Development Economics*, 2008, pp. 3086–3112.

Muralidharan, Karthik and Paul Niehaus, "Experimentation at Scale," *Journal of Economic Perspectives*, 2017, *31* (4), 103–124.

Perla, Jesse and Christopher Tonetti, "Equilibrium Imitation and Growth," *Journal of Political Economy*, 2014, *122* (1), 52–76.

_ **,** _ **, and Michael E. Waugh**, "Equilibrium Technology Diffusion, Trade, and Growth," *American Economic Review*, 2021, *111* (1), 73–128.

Romer, Paul M., "Endogenous Techniological Change," *Journal of Political Economy*, 1990, *98* (5), S71–S102.

Schennach, Susanne M., "Mismeasured and unobserved variables," in Steven N. Durlauf, Lars Peter Hansen, James J. Heckman, and Rosa L. Matzkin, eds., *Handbook of Econometrics*, 2020, pp. 487–565.

Wallskog, Melanie, "The Mechanisms and Implications of Entrepreneurial Spillovers Across Coworkers," 2021. Working Paper.

World Bank, "Project Appraisal: Proposed Credit in the Amount of SDR 62.7 Million (US$86 Million Equivalent) to the Republic of Malawi for a Financial Inclusion and Entrepreneurship Scaling Project," 2020. World Bank Report Number PAD3415.

Table of Contents for Online Appendix

A Additional Details from the Main Text

A.1 RCT: Baseline Profit Distributions

Figure 7: Baseline Profit Distributions

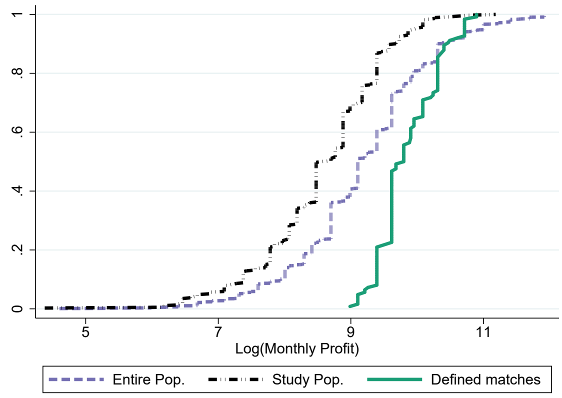

A.2 RCT: Balance Tests

Our basic balance check is

$$y_{i0} = \alpha_0 + \alpha_1 \mathbf{T}_i + \varepsilon_i,$$

where y_{i0} is the baseline outcome for individual i and $T_i = 1$ if i is eventually treated. Those results are in Table 4.

We conduct the second balance test

$$y_{i0} = \alpha_0 + \sum_{j=L,M,H} \alpha_j \mathbf{T}_{ij} + \varepsilon_i$$

where the indicator now depends on whether firm i is a treatment firm matched with a bottom 25th percentile (denoted M_{iL}), 25-75 percentile (T_{iM}), or top 25 percentile firm (T_{iH}) in terms of baseline profitability.[25] Table 5 reports the results. The only significant difference is in age, and the magnitude is small.

A.3 RCT: Empirical Impact on More Productive Member of the Match

Since the more productive members of treatment matches were not randomly selected, we require a different approach to identify any effect on these business owners. Brooks

[25] We have experimented with a number of different ways to compute the balance table, and all show the same results.

Table 4: Balance Test from (from Brooks et al., 2018)

	Control Mean (1)	Mentor - Control (2)
Firm Scale:		
Profit (last month)	10,054	-975.25
		(1186.76)
Firm Age	2.39	-0.05
		(0.23)
Has Employees?	0.21	-0.02
		(0.05)
Number of Emp.	0.21	0.02
		(0.06)
Business Practices:		
Offer credit	0.74	-0.02
		(0.06)
Have bank account	0.30	-0.03
		(0.06)
Taken loan	0.14	-0.05
		(0.04)
Practice accounting	0.11	0.00
		(0.04)
Advertise	0.07	0.04
		(0.03)
Sector:		
Manufacturing	0.04	-0.03
		(0.02)
Retail	0.69	0.03
		(0.06)
Restaurant	0.14	-0.02
		(0.05)
Other services	0.17	0.06
		(0.05)
Owner Characteristics:		
Age	29.1	-0.25
		(0.64)
Secondary Education	0.51	-0.00
		(0.06)

Table Notes: Columns 1-2 are the coefficient estimates from the regression $y_i = \alpha + \beta T_i + \varepsilon_i$, where T_i is an indicator for treatment. Column 1 is $\widehat{\alpha}$. Statistical significance at 0.10, 0.05, and 0.01 is denoted by *, **, and, ***.

et al. (2018) details these results, but Figure 8 reproduces some key results for simplicity's sake. Figure 8 suggests no statistically discernible discontinuity around the cutoff.

We next test this more formally. In particular, letting $\overline{\varepsilon}$ be the cut-off value for mentors, we run the regression

$$\pi_i = \alpha + \tau D_i + f(N_i) + \nu_i \tag{A.1}$$

Table 5: Balancing Test at Baseline

	Control Mean (1)	T_L - Control (2)	T_M - Control (3)	T_H - Control (4)
Firm Scale:				
Profit (last month)	10,054	-732.65	-1337.06	-760.08
		(1314.56)	(1393.38)	(2128.41)
Firm Age	2.39	0.04	-0.19	0.08
		(0.28)	(0.30)	(0.46)
Has Employees?	0.25	-0.10	-0.07	0.10
		(0.07)	(0.07)	(0.11)
Number of Emp.	0.23	-0.05	0.00	0.18
		(0.08)	(0.08)	(0.13)
Business Practices:				
Offer credit	0.74	-0.07	0.04	-0.03
		(0.07)	(0.08)	(0.12)
Have bank account	0.30	-0.04	-0.05	0.06
		(0.07)	(0.08)	(0.12)
Taken loan	0.14	-0.07	-0.06	0.03
		(0.05)	(0.05)	(0.08)
Practice accounting	0.01	-0.01	0.01	-0.01
		(0.01)	(0.02)	(0.02)
Advertise	0.07	0.04	0.01	0.11
		(0.05)	(0.05)	(0.07)
Sector:				
Manufacturing	0.04	-0.02	-0.04	-0.04
		(0.02)	(0.03)	(0.04)
Retail	0.69	-0.03	0.00	-0.10
		(0.08)	(0.08)	(0.12)
Restaurant	0.14	-0.06	0.00	0.03
		(0.05)	(0.06)	(0.09)
Other services	0.17	0.09	0.02	0.07
		(0.06)	(0.07)	(0.10)
Owner Characteristics:				
Age	29.1	0.92	-1.88	0.50
		(0.79)	(0.84)**	(1.28)
Secondary Education	0.51	0.02	-0.08	0.13
		(0.08)	(0.09	(0.13)

Table notes: Columns 1-4 are the coefficient estimates from the regression above, with column one being the estimate of the constant $\widehat{\alpha}_0$. Statistical significance at 0.10, 0.05, and 0.01 is denoted by *, **, and, ***. All constants are significant at one percent.

where π_i is profit, $D_i = 1$ if individual i was chosen as a mentor ($\widehat{\varepsilon}_i \geq \overline{\varepsilon}$, $f(N_i)$ is a flexible function of the normalized running variable $N_i = (\widehat{\varepsilon}_i - \overline{\varepsilon})/\sigma_\varepsilon$, and ν_i is the error term. The parameter τ captures the causal impact of being chosen as a mentor. We use local linear regressions to estimate the treatment effects on profit and inventory, along with business practices of record keeping and marketing. The results are in Table 6, and we find that being a mentor has no statistically significant effect on profits. Moreover, there is no change in marketing or record-keeping practices, which

Figure 8: Profit for mentors and non-mentors (from Brooks et al., 2018)

(a) Full sample

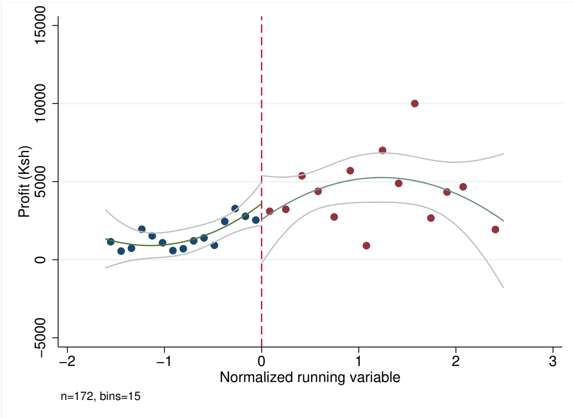

(b) IK Optimal Bandwidth

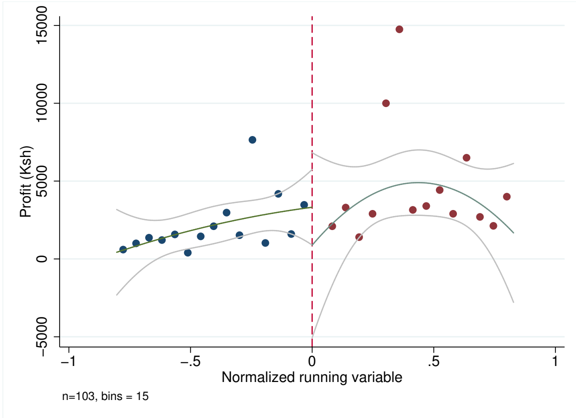

Figure notes: Figure 8 plots profit along with a fitted quadratic and its 95 percent confidence interval. Figure 8a uses the entire sample, while Figure 8b uses the Imbens and Kalyanaraman (2012) procedure to choose the optimal bandwidth. Both use 15 bins on either side of the cutoff.

one might associate with ability. There is some evidence that inventory spending decreases, but it cannot be statistically distinguished from zero. Overall, we find little evidence that entering into a match changes either business scale or business practices for the more productive member of the match. This is consistent with the max function in the forward equation for ability (equation 5.1), which is assumed here and in much of the existing literature.

Table 6: Regression discontinuity results for matched firm treatment effect (Brooks et al., 2018)

Percent of IK	Scale		Practices	
optimal bandwidth	Profit	Inventory	Marketing	Record keeping
100	-503.18	-3105.87	0.01	0.02
	(1321.82)	(2698.11)	(0.11)	(0.18)
150	300.19	-2585.22	0.01	0.07
	(1407.26)	(2291.34)	(0.09)	(0.14)
200	322.09	-123.59	0.01	0.10
	(1324.17)	(1964.08)	(0.08)	(0.13)
Treatment Average	4387.34	8435.79	0.08	0.85
Control Average	1794.09	4039.20	0.13	0.63

Table notes: Statistical significance at 0.10, 0.05, and 0.01 is denoted by *, **, and, ***. Profit and inventory are both trimmed at 1 percent.

A.4 Model: Increasing θ to match RCT treatment effect

In the main text, we increase θ from $\theta = -0.417$ to $\theta = -0.063$. We replicate the results here under a different quantitative experiment: we increase θ such that the policy change induces the same partial equilibrium effect in the model as in our empirical results. Specifically, we create a treatment and control group identical to our empirical RCT. We then shock the treatment group by allowing them to draw from a new source distribution $\widehat{M}(\hat{z};\theta^T)$ instead of $\widehat{M}(\hat{z};\theta)$. We set θ^T such that running the RCT in the model delivers an identical treatment effect to our empirical results.[26] This implies that θ increases to $\theta = 0.512$.

Table 10 replicates Table 3 from the main text under this different quantitative experiment.

Table 7: Equilibrium Moments

	(1) Baseline	(2) Fixed Wage	(3) New Equilibrium
Income	1.00	1.35	1.59
Ability	1.00	1.38	1.67
Aggregate Labor Supply	1.00	0.63	0.90
Wage	1.00	1.00	1.76

Table notes: Column (1) is the initial equilibrium, normalized to one. (2) and (3) report the new stationary equilibrium after shocking the matching technology, where (2) holds the wage fixed at its baseline level and (3) allows it to adjust.

Overall, average income rises by 59 percent. When we decompose the aggregate change, 57 percent of the total ability change comes from the direct effect on ability (Column 2) and the remaining comes from the amplification through the wage. Similarly for income, the direct effect accounts for 59 percent. The mechanisms for these changes are discussed in the main text. Note, however, that these results point to a larger role for price amplification (i.e., the price amplification accounts for 41 percent of the average income gain here and 37 percent in the main text). This occurs because the induced wage change is convex in the imposed θ change. Thus, the results point to a larger role for price amplification as the aggregate policy change becomes larger.

[26] In the empirics, we let treated firms draw from the source distribution $\widehat{H}_T$. The idea is the same, but does not require the treatment source distribution to be in the same functional class as $\widehat{M}$.

B Additional Evidence on Diffusion from Second RCT

B.1 Theory

We begin with a simple, partial equilibrium model that highlights the distinction between two models. The first is a model of "mentor as manager," in the sense that the mentor can replace the mentee's ability in the mentee's business. However, no ability transfers from mentor to mentee. The second is a model of "diffusion," in which the mentor passes ability to the mentee. The model is developed here in a simple, partial equilibrium way to highlight the test we run below.

An individual i is defined by baseline ability $z_{i,t}$ at period t. She then receives a match, whose ability is $\hat{z}_{i,t}$. Both her initial ability and her match's ability allow i's business to be more profitable at t. We define $z^e_{i,t} = g(z_{i,t}, \hat{z}_{i,t})$ as i's effective ability at time t and that $z_{i,t+1} = z^e_{i,t}$. Her profit function is $\pi(z^e_{i,t}, \hat{z}_{i,t})$.

Profit can increase either because i's effective ability increases or because her mentor is better. A better mentor enters into both arguments. The two models are restrictions on derivatives of π and g. The "mentor as manager" model requires $g_2(z, \hat{z}) = 0$ and $\pi_2(z^e, \hat{z}) > 0$. These restrictions state that no knowledge of the match transfers to the agent ($g_2 = 0$), but that the match's ability can directly increase the firm's profitability ($\pi_2 > 0$). On the other hand, the diffusion model requires $g_2(z, \hat{z}) > 0$ and $\pi_2(z^e, \hat{z}) = 0$. These restrictions imply that the mentor can affect profitability by transferring ability to the mentee ($g_2 > 0$) but does not directly affect firm profitability through, for example, managing the firm for the mentee ($\pi_2 = 0$).

With these restrictions, we can derive some predictions that distinguish the two models. First, however, notice that both models imply a positive treatment effect from the random matching program described in the main text as long as $\pi_1 > 0$. With a slight abuse of notation, the realized profit $\tilde{\pi}_{i,t}$ is

$$\frac{\partial \tilde{\pi}_{i,t}}{\partial \hat{z}_{i,t}} = \begin{cases} \pi_2(z^e_{i,t}, \hat{z}_{i,t}) > 0 & \text{if mentor-as-manager} \\ \pi_1(z^e_{i,t}, \hat{z}_{i,t}) g_2(z_{i,t}, \hat{z}_{i,t}) > 0 & \text{if diffusion.} \end{cases}$$

The baseline experiment in the main text cannot directly distinguish these two models.

We now introduce an additional step in the intervention that allows us to distinguish them. Now, after i meets her match, i then goes on to mentor j at $t + 1$. Then we can write j's match as $\hat{z}_{j,t+1} = z^e_{i,t} = g(z_{i,t}, \hat{z}_{i,t})$ and her productive ability at $t + 1$

as $z^e_{j,t+1} = g(z_{j,t+1}, g(z_{i,t}, \hat{z}_{i,t}))$. Applying the definitions of π and g and taking some derivatives yields

$$\frac{\partial \tilde{\pi}_{j,t+1}}{\partial \hat{z}_{i,t}} = \begin{cases} 0 & \text{if mentor-as-manager} \\ \pi_1(z^e_{j,t+1}, \hat{z}_{j,t+1}) g_2(z_{j,t+1}, \hat{z}_{j,t+1}) g_2(z_{i,t}, \hat{z}_{i,t}) > 0 & \text{if diffusion.} \end{cases}$$

This result says something intuitive. If a match only affects contemporaneous profit, then looking a match's previous match should be irrelevant. This is the case of the mentor-as-manager model. On the other hand, if some of i's original match stays with i, then original match quality $\hat{z}_{i,t}$ shows up in j's span of control.

On the other hand, both models predict that j's profit $\pi_{j,t+1}$ should correlate with i's initial ability $z_{i,t}$ because we allow for persistence in profit.

$$\frac{\partial \tilde{\pi}_{i,t}}{\partial z_{i,t}} = \begin{cases} g_1(z_{i,t}, \hat{z}_{i,t}) \pi_2(z^e_{j,t+1}, \hat{z}_{j,t+1}) > 0 & \text{if mentor-as-manager} \\ g_1(z_{i,t}, \hat{z}_{i,t}) \pi_1(z^e_{j,t+1}, \hat{z}_{j,t+1}) g_2(z_{j,t+1}, \hat{z}_{j,t+1}) > 0 & \text{if diffusion.} \end{cases}$$

These results form the basis of our test.

B.2 Empirical Implementation

In part to test this idea, we conducted a second experiment in 2020 that follows the procedure laid out above. We created initial, randomly-paired matches in Dandora in the same way as the baseline experiment discussed in the main text. We then randomly selected mentees to mentor randomly selected control firms after 6 months. Mentees were only made aware of this second round matching program at the 6 month mark. Their previous learning was therefore not influenced by their potential selection into an eventual mentor role.

We refer to this second group of treated firms as the "RCT2 firms." We tracked RCT2 firms for 4 waves for 6 months, with the random matches occurring between the first two waves.

We test between the two models with the regression run only on treated firms in RCT2,

$$\log(\pi_{i,t}) = \tilde{c} + \beta_1 \log\left(\frac{\hat{\pi}^1_{i,0}}{\pi_{i,0}}\right) + \beta_2 \log\left(\frac{\hat{\pi}^2_{i,0}}{\pi_{i,0}}\right) + \gamma_t + \lambda_i + \varepsilon_{i,t}. \tag{B.1}$$

Here $\pi_{i,t}$ is the profit of a firm treated in RCT2, $\hat{\pi}^1_{i,0}$ is the baseline profit of i's mentor and $\hat{\pi}^2_{i,0}$ is the baseline profit of i's mentor's mentor. That is, if k mentors j in RCT1, then j mentors i in RCT2, $\hat{\pi}^1_{i,0} := \pi_{j,0}$ and $\hat{\pi}^2_{i,0} = \pi_{k,0}$. We allow for wave ($\gamma_t$) and individual ($\lambda_i$) fixed effects to test robustness. The baseline values are from the period before k mentors j. For robustness below, we also replace $\log(\hat{\pi}^1_{i,0}/\pi_{i,0})$ with $\log(\max\{1, \hat{\pi}^1_{i,0}/\pi_{i,0}\})$ and similarly for the second term.

The key test of the models is whether or not $\widehat{\beta}_2 > 0$. That is, if our economic environment is better characterized by the mentor replacing the span of control of the mentee for the duration of the match, then $\hat{\pi}^2_{i,0}$ should have no predictive power for i's profit. The results are in Table 8 and we find substantial evidence that the *original* mentor's causes higher profit.

Table 8: Testing Two Models in New RCT

	(1)	(2)	(3)	(4)
Current mentor (β_1)	0.162	0.256	0.336	0.483
	(0.114)	(0.111)**	(0.002)***	(0.130)***
Mentor's mentor (β_2)	0.416	0.489	0.424	0.613
	(0.100)***	(0.104)***	(0.078)***	(0.097)***
R^2	0.468	0.470	0.404	0.619
Average of Outcome Variable	6.710	6.710	6.710	6.710
Include Indiv FE?	N	Y	N	Y

Table notes: Columns 3 and 4 replace the ratios in (B.1) with the max of the profit ratio and 1. Standard errors are in parentheses. Statistical significance at 0.10, 0.05, and 0.01 is denoted by *, **, and, ***.

The first row (β_1) of Table 8 shows that those with better mentors cause higher profit, confirming the main text. The second row (β_2) shows that conditional on that baseline mentor quality, having a mentor *whose own mentor* was higher quality causes higher profitability. This prediction is inconsistent with the "mentor as manager"-style model and instead shows that profit-relevant ability is being transmitted between firms in a match.

C Examples of Different Matching Processes

In the main body of the paper, we provided two examples of matching processes that fall under our assumptions, and we detail additional versions here.

C.1 Noise in the Imitation Process

An agent with ability z receives new arrivals of ideas that have two components: z_m that comes from a random match from another agent, and γ a random innovation on that idea. Then $\hat{z} = \gamma^{1/\theta} z_m$. Here, z_m is a uniform draw from the distribution of productivities. Then if γ has a cumulative density function given by Γ, then:

$$\widehat{M}(c) = Prob(\hat{z} \leq c) = Prob(z_m \leq c\gamma^{-1/\theta}) = \int M(c\gamma^{-1/\theta}) d\Gamma(\gamma) \tag{C.1}$$

C.2 Effort Choice and Bargaining

Each period, every agent characterized by ability z is matched to an agent that owns a potential imitation opportunity z_m as a uniform draw from the distribution of operating firms M. The agent has an effort endowment of 1 that must be divided between imitation and providing a utility benefit to the owner of the imitation opportunity z_m. If $z \geq z_m$, then no effort is put into imitation and $\hat{z} = z$. If $z_m > z$, then the agent and the owner of the imitation opportunity must first agree on the distribution of effort, then the choice of effort x and the values of z and z_m together generate the value of $\hat{z}$ for the agent in that period according to:

$$\hat{z} = \left(\frac{z_m}{z}\right)^x z \tag{C.2}$$

That is, by putting in more effort $x \in [0, 1]$ the agent is able to close the gap between their z and z_m. The benefit to the owner of z_m is given by the function $b(x)$, which is decreasing in x.

Agents and owners of imitation opportunities have one-off interactions and each receive 0 benefit if no agreement is made. They bargain over the assignment of the agent's effort between imitation and utility benefits for the owner of the imitation opportunity according to a Nash bargaining problem where the bargaining weight of the agent is θ. The bargaining problem is:

$$\max_{x\in[0,1]} \left(\left[\frac{z_m}{z}\right]^x z\right)^\theta b(x)^{1-\theta} \tag{C.3}$$

Suppose that $b(x)$ is given by $b(x) = 1 - x$. Then it is easy to show that:

$$x = \max\left[0, 1 - \frac{1-\theta}{\theta \log(z_m/z)}\right] \tag{C.4}$$

$$\hat{z} = \max\left[z, z_m e^{1-1/\theta}\right] \tag{C.5}$$

As expected, the more bargaining power that the learning agents have, the greater is x, resulting in greater $\hat{z}$.

Note that, in the model, draws of imitation opportunities $\hat{z} < z$ are not useful. Hence, the distribution $\widehat{M}$ can be written, for any value c, as:

$$\widehat{M}(c) = Prob(\hat{z} \leq c) = Prob(z_m e^{1-1/\theta} \leq c) = Prob(z_m \leq ce^{1/\theta-1}) = M(ce^{1/\theta-1}) \tag{C.6}$$

or following the notation more standard in the paper:

$$\forall z, \widehat{M}(\hat{z}; z, \theta) = M(\hat{z}e^{1/\theta-1}) \tag{C.7}$$

C.3 Deterministic Assignment

Here we consider a case where $\widehat{M}$ arises when all agents can interact with one another and sort into relationships endogenously. Suppose that every agent with ability $\hat{z}$ has the option to influence any other agent that has ability z. Every agent can only be influenced by one other agent each period, and they always prefer to be influenced by the highest ability possible.

The utility of an agent with ability $\hat{z}$ influencing an agent with ability z is given by:

$$\frac{\hat{z}}{z} - 1 - \frac{1}{2\theta}\left(\frac{\hat{z}}{z} - 1\right)^2 \tag{C.8}$$

That is, the agent with $\hat{z}$ gains benefit in proportion to how large the benefit is for the other agent, but their cost is quadratic in the distance between their productivities. For example, the influencer is happy when the other agent is helped by their influence, but it takes more effort to influence when the distance between them is great. Therefore, if there is a continuous distribution of $z < \hat{z}$, the ideal agent that the influencer would like to interact with has ability:

$$z^*(\hat{z}) = \hat{z}/(1+\theta) \tag{C.9}$$

That is, the lower is the cost of influencing low ability firms, the deeper into the left tail of the distribution is the agent willing to go.

However, since every agent can only be influenced by one agent each period and they strictly prefer to be influenced by agents of higher ability, it is possible that (even if the distribution is continuous) that the ideal agent for $\hat{z}$ is already matched to another influencer. Therefore, intuitively, the probability distribution over assignment between $\hat{z}$ and z is constructed by starting at the upper support of the distribution M, allowing the highest ability firms to choose their most preferred matches, then descending down through the distribution letting each firm choose to influence its preferred firm among those remaining. Note that not all firms need have another firm to influence if their utility from doing so be negative.

Formally, the probability distribution over imitation opportunities can be constructed in the discretized case as follows, when the ability grid takes values $z \in \{z_1, ... z_N\}$, which are ordered ($i < j \implies z_i < z_j$).

Define $\tilde{\mu}(z, \hat{z})$ as the measure of $\hat{z}$ influencing z (a $N \times N$ matrix). We can construct $\tilde{\mu}$ in the following steps given the measure μ of agents of each z type:

1. Let $U(z, \hat{z})$ be the $N \times N$ matrix of utilities of $\hat{z}$ influencing z, and $\tilde{\mu}$ be a $N \times N$ matrix of zeros. Let $\bar{\mu}$ be the $N \times 1$ vector of unassigned influencers and μ_u be the $N \times 1$ vector of unassigned imitators. Set $\bar{\mu} = \mu_u = \mu$, $n = N$, and $m = 1$.

2. Let l be the m-argmax of $U(\cdot, z_n)$. If $U(z_l, z_n) \leq 0$, set $\tilde{\mu}(z_1, z_n) = \mu_u(z_n)$ and skip to step 5.

3. If $\bar{\mu}(z_n) \leq \mu_u(z_l)$, then $\bar{\mu}(z_n) = 0$, $\mu_u(z_l) = \mu_u(z_l) - \bar{\mu}(z_n)$, and $\tilde{\mu}(z_l, z_n) = \bar{\mu}(z_n)$. Skip to step 5. Otherwise, go to 4.

4. If $\bar{\mu}(z_n) > \mu_u(z_l)$, then set $\tilde{\mu}(z_l, z_n) = \mu_u(z_l)$, $\mu_u(z_n) = 0$ and $\bar{\mu}(z_n) = \bar{\mu}(z_n) - \mu_u(z_l)$. Set $m = m + 1$ and return to step 2.

5. Set $n = n - 1$ and $m = 1$. If $n = 0$, go to step 6. Otherwise, go to step 2.

6. Set $\tilde{\mu}(\cdot, z_1) = \tilde{\mu}(\cdot, z_1) + \mu_u$, and stop.

Given this matrix $\tilde{\mu}(z,\hat{z})$, the measure of assignments $\widehat{M}$ is given by:

$$\widehat{M}(\hat{z}_i, z_j) = \frac{\sum_{k=1}^{i} \tilde{\mu}(z_j, \hat{z}_k)}{\mu(z_j)}$$

C.4 Cost to Receive a Match

Firms pay a cost to receive a uniform random draw from the productivity distribution. We model this as a function $f(s;\theta)$, in that a firm that pays cost $f(s;\theta)$ receives a uniform random draw with probability s. We assume that $f_\theta > 0$, so that the cost to achieve any level of s is increasing in θ. In a stationary equilibrium, under standard conditions this implies a stationary decision rule $s(z, M^*;\theta)$ with $s_\theta < 0$. We can write the distribution of draws as

$$\begin{aligned} \widehat{M}(c;z,\theta) &= (1 - s(z, M^*;\theta)) + s(z, mM^*;\theta)M^*(c) \\ &\equiv Q(c;z,M^*,\theta) \end{aligned}$$

Thus, in the stationary equilibrium of this economy, our same procedure goes through. This example provides important context for our set of assumptions – they need not be assumptions only on the primitives of the model. Additional assumptions, such as stationarity, may guarantee the model is covered under our assumptions. The use of stationarity here is similar to its role in the identification procedure of Jarosch et al. (2021), who study learning among German co-workers.

D Extensions of Diffusion Parameter Identification

In this Appendix, we focus on theoretical extensions of the main estimation procedure in the paper to show that the procedure itself is robust to any number of extensions. In Appendix F, we provide a quantitative evaluation of a particular extension related to mis-measurement.

D.1 Semi-parametric identification

Assumption 1 laid out a function form for the law of motion of ability: $z'(z, \varepsilon, \hat{z}) = e^{c+\varepsilon} z^{\rho} \max\left\{1, \frac{\hat{z}}{z}\right\}^{\beta}$. We broaden this in Assumption 5 by replacing the max function,

Assumption 5. *Given ability z this period, an imitation opportunity $\hat{z}$, and a random shock ε, ability next period z' is given by*

$$z'(z, \varepsilon, \hat{z}) = e^{c+\varepsilon} z^{\rho} f\left(\frac{\hat{z}}{z}\right) \tag{D.1}$$

Assuming $f(\hat{z}/z) = \max\{1, (\hat{z}/z)\}^{\beta}$ gives us the original Assumption 1. Proposition 4 summarizes that we can instead estimate the function f using the same data-generating process as in the main text.

Proposition 4. *The data-generating process of Assumption 4 identifies (ρ, f) in equation (D.1) (while maintaining Assumptions 2 and 3 in the main text) by estimating the regression*

$$\log(\pi_i') = \tilde{c} + \rho \log(\pi_i) + f\left(\frac{\hat{\pi}_i}{\pi_i}\right) + \varepsilon$$

This result follows directly from an established literature on partially linear regressions. See, for instance, Yatchew (1997) on differencing estimators in this class of models.[27] Härdle et al. (2000) provides a detailed review.

D.2 More general relationship between observables and z

In Assumption 2 we assume that some observable (e.g., profit) π has the characteristic that $\pi \propto z$. We extend that here.

[27] The basic idea is that by ordering the data such that $\hat{\pi}_1/\pi_1 < \hat{\pi}_2/\pi_2 \ldots < \hat{\pi}_N/\pi_N$, once can difference out the nonlinear f in the limit under conditions that guarantee that the gap between i and $i+1$ goes to zero. This allows straightforward OLS to estimate ρ. Then f can be estimated non-parametrically by any number of methods.

Assumption 6. *There exists a known function $g : \mathbf{X} \to \mathbb{R}_{++}$ that maps observable characteristics $\mathbf{x}$ to ability z up to a potentially unknown constant of proportionality. That is, $g(\mathbf{x}) = Cz$ for some potentially unknown constant C.*

By assuming $\pi \in \mathbf{x}$ and $g(\mathbf{x}) = \pi$, we recover the original assumption $\pi \propto z$. But Assumption 6 allows for more complicated possibilities. For example, one could estimate a production function using the control group panel data. Such a procedure would imply a mapping $g(y, n, k) = Cz$. That is, it takes output data y and input data for labor and capital (n, k) (or any other input bundle) and infers the value z.[28]

Proposition 5. *The parameters (β, ρ, θ) are identified when we replace Assumption 2 with Assumption 6.*

This follows almost directly, as the regression

$$\log(g(\mathbf{x}')) = c + \rho \log(g(\mathbf{x})) + \beta \log\left(\max\left\{1, \frac{g(\hat{\mathbf{x}})}{g(\mathbf{x})}\right\}\right) + \varepsilon,$$

is straightforwardly estimated with known g. By Assumption 6 this collapses to the required equation that gives (β, ρ). The second step goes through with the same adjustment. The key to our procedure is not the proportionality of any one variable with z, but a proportional mapping between *any* set of observables and z.

D.3 Mis-measurement

We now assume that profit is mis-measured. This affects our estimation procedure, introducing bias into the parameters. The regression error is not additively separable from the true value in our non-linear model, which implies that standard instrumental variable methods to correct linear measurement error no longer hold. Yet, there is a substantial and active literature on mis-measurement in non-linear models that provides a number of ways to overcome this issue.

We discuss this issue in the context of a more general ability law of motion, to emphasize that it does not depend on the specific choices we have made on functional forms. We provide a quantitative evaluation of these issues in Appendix F.

[28] g is not the production function here, but is inferred from it.

Assumption 7. *The law of motion for diffusion can be written as*

$$log(\pi') = \sum_{j=1}^{M} \beta_j g_j(\vec{\pi}) + \varepsilon$$

where $\vec{\pi} = (\pi, \hat{\pi})$ *for a known function* $(g_j)_{j=1}^{M}$.

Note that we have already imposed the maintained ability to move between ability z and profit π. The key additional assumption is an adjustment to Assumption 4, which governs the type of data to which we have access.

Assumption 8. *Profit is measured with error, and we observe two outcomes that are mis-measured versions of the true value,* $\vec{\pi}^* = (\pi^*, \hat{\pi}^*)$. *We denote* $\vec{\pi}^k$ *as the two entries in the vector. We denote these observable values as* $(\pi_1^k, \pi_2^k), k = 1, 2$. *The measurement error is classical, so that for each individual* i *we observe*

$$\begin{aligned} \vec{\pi}_{1i}^k &= \vec{\pi}_i^{*k} + \nu_{1i}^k, \quad k = 1, 2 \\ \vec{\pi}_{2i}^k &= \vec{\pi}_i^{*k} + \nu_{2i}^k, \quad k = 1, 2 \end{aligned}$$

where ν_1 *and* ν_2 *are unobserved disturbances. We assume the following relationships between the measurement error and true values:*

$$\mathbb{E}[\nu_1^k | \pi^{*k}, \nu_2^k] = 0, \quad k = 1, 2$$
$$\nu_2^k \textit{ is independent from } \vec{\pi}^*, \nu_2^{-k}, \textit{ where } -k \neq k$$

The assumption of a repeated measurement opens up a suite of tools related to repeated measurement adjustments in non-linear estimation. While other methods exist to solve this problem, the existence of this approach has the benefit of almost always being available in a firm-level survey.[29]

The basic idea behind this identification approach comes from Kotlarski's Lemma, which in $\mathbb{R}^1$ is

$$\phi_{\pi^*}(t) = \exp\left(\int_0^t \frac{\mathbb{E}[i\pi_1 e^{it\pi_2}]}{\mathbb{E}[e^{it\pi_2}]} \right)$$

and $\phi_{\pi^*}(t)$ is the characteristic function $\phi_{\pi^*}(t) = \int_{\mathbb{R}} e^{it\pi^*} f_\pi^*(\pi^*)dx$. An inverse Fourier transform gives us the distribution of true values f_{π^*}, which can be used to construct

[29] For example, π_1 could be profit asked directly while π_2 could be measured as revenue minus costs. Another would be to use the same variable measured at two points in time, as highlighted by Schennach (2020). Finally, many models allow relationships between input expenditures, revenue, and profit that could be utilized, albeit with more structure than we have here. An example is a Cobb-Douglas production function with competitive factor markets.

the relevant estimator moments.

Our model requires $\vec{\pi}^* = (\pi^*, \hat{\pi}^*) \in \mathbb{R}^2$, which introduces some complications. The second part of Assumption 8 provides necessary conditions for identification in a multidimensional nonlinear model.

Proposition 6. *If $\mathbb{E}[|\vec{\pi}^k|]$ and $\mathbb{E}[|\eta_1^k|]$ are finite, then there exists a closed form for any function $\mathbb{E}[u(\vec{\pi}^*, \beta)]$ whenever it exists.*

Proof. Provided in Schennach (2004). ■

Schennach (2004) provides the details of the approach and how to develop an estimator from Proposition 6. The key feature, however, is that this result allows a broad class of extremum estimators to be deployed to identify β.

D.4 Additional characteristics

One might also suspect that alternative characteristics influence learning. For example, firm owners may retain more ability when meeting with another owner of a similar age. This amounts to allowing β to depend on a set of characteristics of the firm owner $\mathbf{x}$ and her match $\hat{\mathbf{x}}$. In this case, we can write

$$\log(\pi_i') = c + \rho \log(\pi_i) + \beta(\mathbf{x}, \hat{\mathbf{x}}) \log\left(\max\left\{1, \frac{\hat{\pi}_i}{\pi_i}\right\}\right)$$

or, binning characteristics $\mathbf{X} \times \hat{\mathbf{X}}$ in some way,

$$\log(\pi_{ib}') = c + \rho \log(\pi_{ib}) + \sum_{b=1}^{B} \beta_b \log\left(\max\left\{1, \frac{\hat{\pi}_{ib}}{\pi_{ib}}\right\}\right)$$

This regression identifies $(\rho, \beta_1, \ldots, \beta_B)$ under the same assumptions as in the main text. The second step follows with a slight adjustment to Assumption 3. We assume there is a discrete distribution over types Γ_b and, with a slight abuse of notation, re-write the draws over types and profit as

$$\widehat{M}(\hat{z}, b; z, \theta) = \widehat{M}_b(\hat{z}; z, \theta)\Gamma_b$$

As long as $\widehat{M}_b$ has the same properties as Assumption 3, the second step of the procedure similarly identifies θ.

E Characterizing the Efficient Allocation from the Social Planner's Problem

Here, we lay out the solution to the social planner's problem, and show that it similarly relies on properly measuring the intensive and extensive margins.

Before doing so, one conceptual issue to deal with is the role of suppliers. We exclude them from our measure of welfare. In our context, these suppliers take their profits out of Dandora, so that buying intermediates does indeed involve resources exiting the economy. To operationalize this idea, we assume that the price paid follows the same solution as the Nash bargaining problem solved in Proposition 1, in that the price p_x is given

$$p_x = \left(\frac{1-\eta-\nu(1-\eta-\alpha)}{\alpha} \right) e^{-s} z^{\frac{\alpha+\eta-1}{\alpha}}.$$

We will write $p_x(z,s)$ to denote this price.[30]

With those details, we are ready to define the planner's problem. The social planner allocates occupations o, firm inputs (x,n) and supplier search intensity s for each agent to maximize *ex ante* utility, subject to the relevant aggregate resource constraints. Since we will focus on the stationary equilibrium, we drop the dependence of the decision rules on the aggregate state M for some notational simplicity. Defining the value to an agent with ability z as

$$\widetilde{v}(z) = \omega \log(c(z)) + (1-\omega)\log(1-s(z)) + (1-\delta) \int_{\varepsilon} \int_{\hat{z}} \widetilde{v} \left(e^{c+\varepsilon} z^{\rho} \max\left\{1, \frac{\hat{z}}{z}\right\}^{\beta} \right) d\widehat{M}(\hat{z}) dF(\varepsilon),$$

we can write the planner's problem recursively as

$$\max_{o(\cdot),c(\cdot),x(\cdot),n(\cdot),s(\cdot)} \int_0^{\infty} \widetilde{v}(z) dM(z) \tag{E.1}$$

[30] There is nothing conceptually difficult about including these suppliers in the measure of welfare. Our goal is only to remain faithful to the economic environment from which the empirics are derived. A further benefit of this assumption is that it focuses attention on the role of the diffusion externality, as opposed to other types of inefficiencies that arise from the bargaining protocol.

$$s.t. \quad \int_{o(z)=1} \Big(x(z)^{\alpha} n(z)^{\eta} - p_x(z, s(z)) x(z) \Big) dM(z) = \int_0^{\infty} c(z) dM(z) \tag{E.2}$$

$$\int_{o(z)=1} n(z) dM(z) = \int_{o(z)=0} dM(z) \tag{E.3}$$

$$M(z') = \delta G(z') + \tag{E.4}$$

$$(1-\delta) \int_0^{\infty} \int_0^{\infty} F(\log(z') - \rho \log(z) - \beta \log(\max\{1, \hat{z}/z\}) - c) d\widehat{M}(\hat{z}) dM(z)$$

$$\widehat{M}(\hat{z}; M) = \left(\frac{\int_0^{\hat{z}} \phi(z, M) dM(z)}{\int_0^{\infty} \phi(z, M) dM(z)} \right)^{\frac{1}{1-\theta}}. \tag{E.5}$$

While complicated looking, this problem has a straightforward interpretation. The planner's objective is to maximize the expected value of $\widetilde{v}$. The first two constraints are the aggregate resource constraints – (E.2) determines the resources that can be allocated to consumption, while (E.3) equalizes labor supply and demand. The latter two constraints show that the planner internalizes how her decisions affect the evolution of the aggregate state (via E.4) and the imitation opportunities that arise from it (via E.5). These constraints highlight the planner's ability to overcome the diffusion externality, in that she takes into account the implications of occupational choice on learning opportunities in a way that individual agents do not.

We measure the difference in welfare between the allocation chosen by the planner and the baseline *laissez faire* equilibrium in consumption-equivalent terms. That is, we ask by what percentage we would we have to increase each agent's consumption in every state and time period to equalize average utility between the two economies. This difference is our measure of the aggregate importance of diffusion.

While the planner's problem does not admit a closed-form solution, we can derive some implications for comparison to the baseline *laissez faire* equilibrium.

Proposition 7. *The planner chooses a cutoff rule for occupations, $\underline{z}$, such that all $z \geq \underline{z}$ operate firms. Consumption c_p is constant across agents and given by the constant returns to scale aggregate production function $c_p = A N_s^{\frac{\eta}{1-\alpha}} Z^{\frac{1-\alpha-\eta}{1-\alpha}}$, where A is a constant and the two aggregate inputs are*

$$N_s = \int_0^{\underline{z}} dM(z) \qquad , \qquad Z = \int_{\underline{z}}^{\infty} z \exp(s(z))^{\frac{\alpha}{1-\alpha-\eta}} dM(z).$$

Moreover, if $s(\underline{z}) > 0$ (which we verify at our estimated parameters), $s(\cdot)$ is a strictly increasing and concave function of the form

$$s(z) = \left(\frac{1-\alpha-\eta}{\alpha}\right) W_0\left[\left(\frac{\alpha}{\eta+\alpha-1}\right)\exp\left(\frac{\alpha}{\eta+\alpha-1}\right) q(z/\underline{z}, s(\underline{z}))\right] + 1 \quad \forall z \geq \underline{z},$$

where $W_0(\cdot)$ is the principal branch of the Lambert W function and $q(\cdot,\cdot)$ is a function of relative ability $z/\underline{z}$ and supplier search at the cut-off value, $s(\underline{z})$.

Proof. The proof is at the end of the Appendix, in Appendix Section H. ■

One can see the goals of the planner in Proposition 7. Like in the baseline, the planner uses a cut-off rule to define occupations. As we show below, she chooses fewer firms than the baseline equilibrium, a function of the diffusion externality in the baseline economy. Furthermore, she shifts the search for suppliers away from relatively low ability agents. Instead, she takes advantage of complementarity between ability and supplier search effort. This allows higher ability agents to procure resources that can be redistributed to all agents. These incentives are naturally absent in the baseline equilibrium, where individuals consume their income.

In terms of our inability to push further theoretically, the properties of W_0 preclude a closed form solution.[31] We derive these results in Appendix H.4, and proceed with quantitative results in the next section.

E.1 Quantitative Implications

Our interest here is understanding the importance of separately identifying the intensive and extensive margin parameters to understand the welfare gains in the social planner's problem.

We follow a similar (but slightly simpler) approach to the main text. We exogenously vary β, then re-estimate θ to continually match the same average treatment effect. Thus, the average treatment effect is still used as an estimating moment, but our first stage regression (run within the treatment group) is not.

Throughout, we hold the persistence ρ fixed for simplicity. Figure 9 plots the implied value of θ and the consumption-equivalent welfare gains, the latter of which

[31] The Lambert W function is the inverse of $F : x \mapsto x\exp(x)$, and we can restrict attention to the principal branch. The main issue for our purposes is that the only analytical characterization of W_0 is the power series $W_0(a) = \sum_{n=1}^{\infty}((-n)^{n-1}/n!)a^n$, which is of little help in attempting to attain a closed form for the planner's problem here. Put slightly more technically, the supplier search function s is the solution to a delay differential equation in z, but this solution prevents an analytic characterization of the initial condition $s(\underline{z})$.

is normalized to one at our estimated value.

Figure 9: Importance of Separately Identifying Two Diffusion Effects

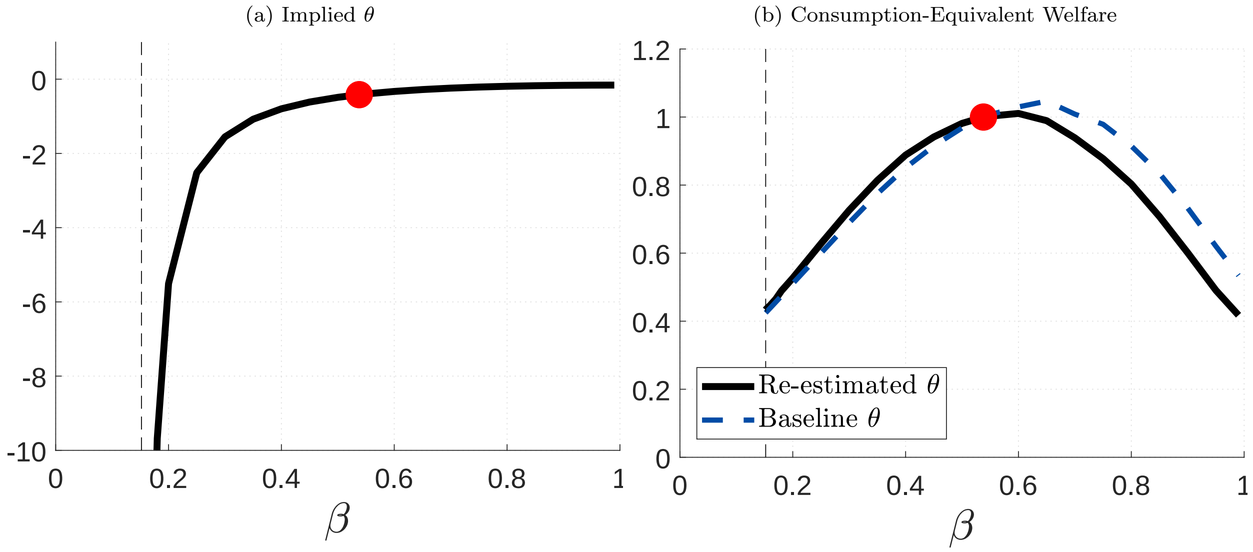

Figure notes: The left panel shows the implied θ required to estimate the same treatment effect observed in the data at the exogenously given β. The dashed vertical line the minimum value of β for which there exists a θ that can rationalize the observed average treatment effect. This is related to the discussion of Γ^{min} defined in Proposition 2. $\theta \to -\infty$ as β approaches this value from above. The right panel shows the implied consumption-equivalent welfare, normalized to one at our estimated parameters. Our estimated values are denoted by the circle in each graph.

The welfare implications are in Figure 9b. We focus first on the solid line. This is our baseline counterfactual in which θ is re-estimated to achieve the same average treatment effect. The differences here can be substantial. Increasing β from 0.25 to to our estimated value 0.538, for example, increases the consumption equivalent welfare gain by 59 percent. Perhaps more surprisingly, overestimating the intensive margin forces can similarly bias downward welfare gains. Assuming $\beta = 0.90$ instead of 0.538 lowers the implied welfare gains by 40 percent. We detail to the economic forces governing these trade-offs in the next section. For now, however, the results show that understanding effecient welfare gains depends critically on separately identifying the parameters governing the extensive (here, θ) and intensive (here, β) margins, as our procedure does.

E.1.1 Understanding Forces Behind This Pattern

To better understand the forces behind the pattern above, we also re-estimate the welfare gains while holding θ fixed at its baseline value. This is the dashed line in Figure 9b. The limited difference between the two lines shows that the welfare gains are primarily driven by the direct effect of β, and not the implied differences in θ.[32] Thus, understanding the main results above primarily requires understanding the impact of changing β. We study that here.

A useful starting point is to re-write total resources in the planner's allocation in terms of semi-elasticities. Recalling the aggregate production function defined in Proposition 7, we can then define the semi-elasticity of consumption with respect to the occupational cut-off, $\varepsilon_{c_p} := (\partial c_p/\partial \underline{z})/c_p$, as the sum of the aggregate input elasticities,

$$\varepsilon_{c_p} = \left(\frac{\eta}{1-\alpha}\right)\varepsilon_{N_s} + \left(\frac{1-\alpha-\eta}{1-\alpha}\right)\varepsilon_Z. \tag{E.6}$$

Equation (E.6) tells us the consumption change induced when the planner slightly shifts the occupational cut-off. There are two components to the consumption gains. The first is static – holding the ability distribution fixed, increasing the cut-off mechanically increases labor supply. This positively affects ε_{N_s} and negatively affects ε_Z. But key in this model is that shifting $\underline{z}$ also affects the ability distribution M through diffusion. This is the dynamic effect on production, as the mass of the population in each occupation changes as M changes.[33] Figure 10 plots the semi-elasticity ε_{c_p}, along with those of the two aggregate inputs defined in (E.6). They are evaluated at the baseline equilibrium cut-off.

The first thing to note about Figure 10a is that the elasticity is positive for any value of β. Thus, no matter the parameters, the planner can increase consumption by transition some baseline entrepreneurs to wage work. This is entirely a function of the diffusion externality. Moreover, as expected, it follows a similar shape to the overall welfare gains. Finally, Figure 10b shows that this is in large part driven by ε_Z. That is, the consumption response is primarily driven by the changes to ability in the economy.

Why, then, does ε_Z take this shape? This is at the heart of understanding how the overall welfare gains vary with the critical parameter β. Figure 11a provides a mechanical rationale: the numerator of ε_Z is approximately linear while the denominator is substantially more convex. But both of these curves have natural economic interpretations. The numerator, $\partial Z/\partial \underline{z}$, measures the change in economy-wide ability

[32] The difference between the two is driven by the fact that the welfare gains are declining in θ.

[33] It is straightforward to show that each of these semi-elasticities can be decomposed into a static and dynamic term, as $\varepsilon_j = \varepsilon_j^S + \varepsilon_j^D$ for $j = N_s, Z$. The static reallocation across occupations plays almost no quantitative role here. Thus, when interpreting the results here, one should think of them as driven by the dynamic diffusion effects.

Figure 10: The Semi-Elasticity of Consumption for the Social Planner

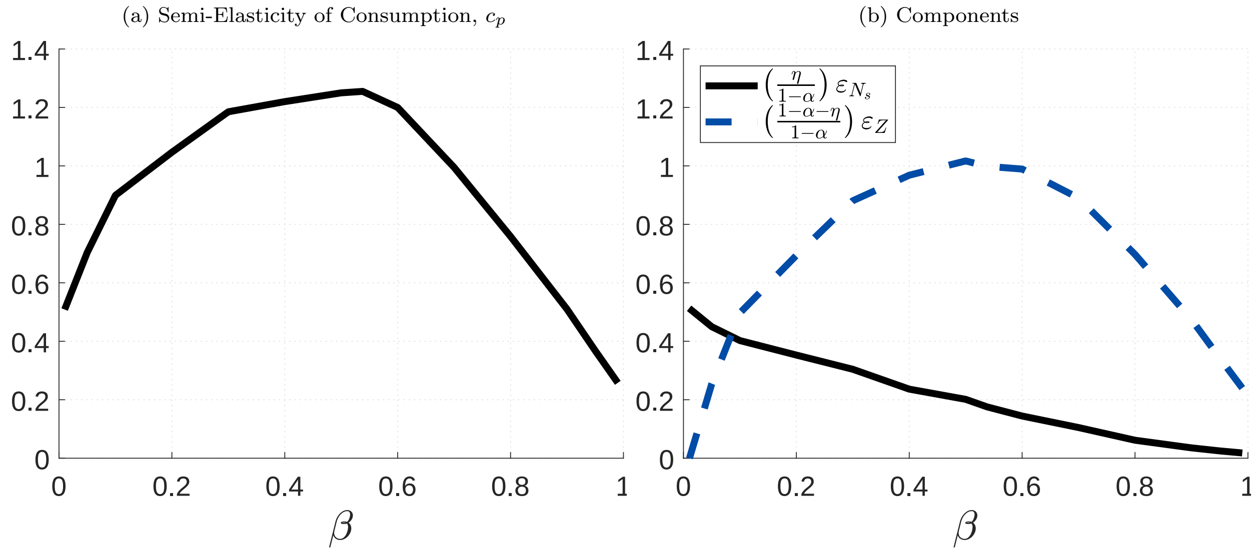

Figure notes: These semi-elasticities are evaluated at the baseline equilibrium cut-off $\underline{z}^*$ and equilibrium distribution $M^*(z;\underline{z}^*)$.

from a marginal increase in the cut-off. It is monotonically increasing – the higher β, the larger the potential impact on ability. This means that all else equal, a higher β induces a larger increase in consumption for the planner.

But the welfare *gains* depend on how that composite ability compares to the baseline equilibrium. This baseline ability is given by Z in Figure 11a, and also depends on β. In particular, higher β increases the gains from a good match. This increases ability and thus drives up the return to labor, which increases the wage. The higher wage induces more wage workers, as Proposition 1 shows that the baseline cut-off has the feature $\underline{z} \propto w^{\frac{1-\alpha}{1-\eta-\alpha}}$ (Figure 11b shows it graphically). These same steps are then compounded in equilibrium. With fewer, more able firms, diffusion accelerates even faster as agents further increase ability. The stationary equilibrium wage then takes the convex share of Figure 11a.

Together, these results highlight the two competing forces in the model, both of which can be seen in Figure 11. At low levels of β, the matching technology limits the aggregate welfare gains. At high levels of β, we see similar welfare gains, but for a different reason. Here, standard equilibrium forces already accomplish most of what the planner would want to accomplish. While the baseline economy is clearly richer, the returns to *additional* intervention by the planner at high β are low. We summarize these forces in Figure 11c, which plots average consumption across the baseline and efficient allocations. Consistent with these results, average baseline consumption grows more slowly than the efficient consumption at low β, then more quickly at high β.

Thus, these two competing economic forces – the race between technology and equilibrium prices in generating diffusion – are critical in generating the non-monotonicity observed in the headline results. Moreover, they are governed are the diffusion parameters we estimate. Thus, estimating these parameters plays an important role in understanding the aggregate implications in the economy.

Much like our main results, the planner's problem shows that separating the extensive and intensive margin is critical to understanding the gains from policy, whether they be a change to the matching technology (as in the main text) or the fully efficient planner's allocation (here).

Figure 11: Equilibrium Forces and the Pattern of ε_Z

(a) Decomposing ε_Z

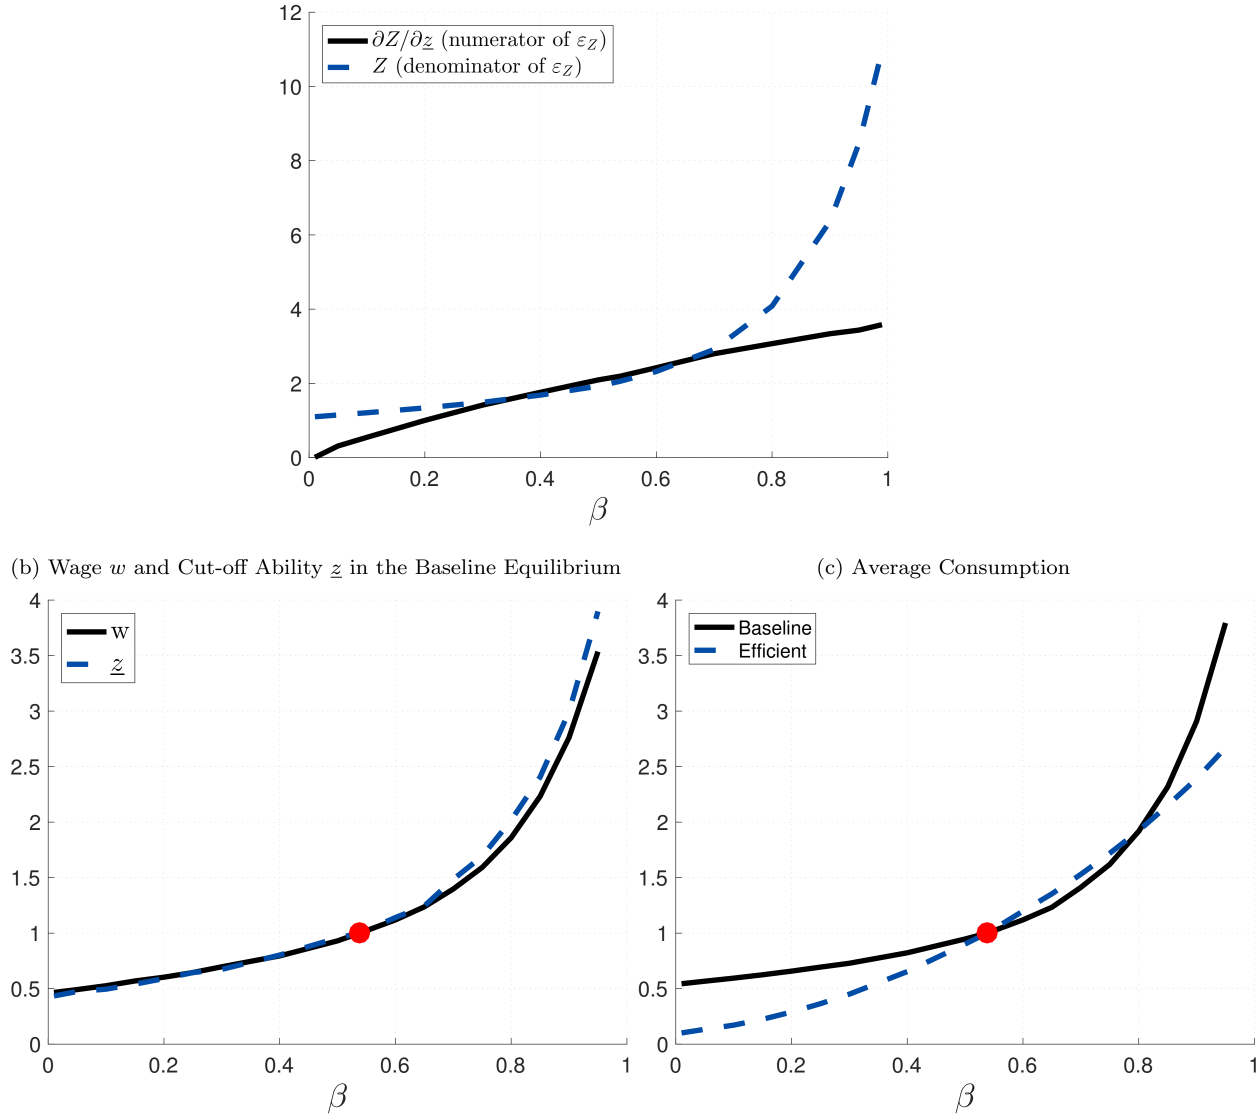

(b) Wage w and Cut-off Ability $\underline{z}$ in the Baseline Equilibrium

(c) Average Consumption

Figure notes: In Figures 11b and 11c, our estimated values are denoted by the circle in each graph and normalized to one.

F Quantitative Implications of Mismeasurement

We explore the quantitative consequences of mis-measuring profit here. We make the following assumption: for all individuals, we observe $\pi = \tau\pi^*$, where $\tau \sim N(0, \sigma_\tau)$ is classical measurement error (in logs), π is observed profit, and π^* is true profit. We assume that $\tau \sim N(0, \sigma_\tau)$, where σ_τ is known but the individual realizations are not. As discussed in Appendix D this can be extended to estimate the distribution of τ. We leave this aside for simplicity here. Define the function

$$g(\pi^*, \hat{\pi}^*; \mathbf{\Gamma}) = c + \rho \log(\pi^*) + \beta \log \left(\max \left\{ 1, \frac{\hat{\pi}^*}{\pi^*} \right\} \right)$$

where $\mathbf{\Gamma} := (c, \rho, \beta)$ is the parameter set. Our first stage regression is $\log(\pi'^*) = g(\pi^*, \hat{\pi}^*; \mathbf{\Gamma})$, but is complicated by the mis-measured profit on the right hand side of this equation. Here, we re-estimate this regression with mis-measured profit and show how the results change.

F.1 Estimation Procedure

Some notation is required. For any variable x, define $\tilde{x} = \log(x)$, $f_x(x)$ as the probability density function, and $\phi_x(t) = \int_{\mathbb{R}} e^{itx} f_x(x) dx$ as its characteristic function.

We first estimate the characteristic functions of the observed π and $\hat{\pi}$,

$$\begin{aligned}
\hat{\phi}_{\tilde{\pi}}(t) &= \left(\frac{1}{n} \sum_{j=1}^{n} e^{it \log(\pi_j)} \right) \phi_{k,\pi}(h_\pi t) \\
\hat{\phi}_{\tilde{\hat{\pi}}}(t) &= \left(\frac{1}{n} \sum_{j=1}^{n} e^{it \log(\hat{\pi}_j)} \right) \phi_{k,\hat{\pi}}(h_{\hat{\pi}} t)
\end{aligned}$$

The first term in parenthesis is the empirical characteristic function using mis-mesaured variables $(\pi_j, \hat{\pi}_j)$. The latter term, $\phi_k(ht)$, is the Fourier transform of a kernel density estimator with bandwidth h.[34]

Since $\phi_{\tilde{\pi}^*}(t) = \hat{\phi}_{\tilde{\pi}}(t)/\phi_{\tilde{\tau}}(t)$ and similarly for $\hat{\pi}$ (due to the independence of π and

[34] A well-known issue with this type of estimation is the inaccuracy of the empirical characteristic function in the tails of the distribution. A kernel density estimate is one version of what is generally referred to as a dampening factor to improve accuracy. The fact that ϕ_k enters multiplicatively follows because a kernel estimator is also a type of convolution.

$\hat{\pi}$), we recover the estimated densities of π and $\hat{\pi}$ from an inverse Fourier transform,

$$\begin{aligned} f_{\pi^*}(\pi^*) &= \frac{1}{2\pi}\int \hat{\phi}_{\pi^*}(t)e^{-it\pi^*}dt \\ f_{\hat{\pi}^*}(\hat{\pi}^*) &= \frac{1}{2\pi}\int \hat{\phi}_{\hat{\pi}^*}(t)e^{-it\hat{\pi}^*}dt \end{aligned}$$

where π in the ratio $1/(2\pi)$ should be understood to be $\pi \approx 3.14$ instead of profit (as to not introduce any additional notation).

With the true distribution functions, we can estimate our regression in any number of ways. We use the minimum distance estimator proposed by Hsiao (1989). That is, we choose parameters $\mathbf{\Gamma} := (\tilde{c}, \rho, \beta)$ to solve

$$\min_{\mathbf{\Gamma}} \sum_{i=1}^{n} \left(\pi_i' - G(\pi_i, \hat{\pi}_i; \mathbf{\Gamma})\right)^2 \tag{F.1}$$

where

$$G(\pi, \hat{\pi}; \mathbf{\Gamma}) = \int\int g(\pi^*, \hat{\pi}^*) f_{\pi^*|\pi}(\pi^*|\pi, \mathbf{\Gamma}) f_{\hat{\pi}^*|\hat{\pi}}(\hat{\pi}^*|\hat{\pi}, \mathbf{\Gamma}) d\pi^* d\hat{\pi}^*$$

The minimum distance estimator in (F.1) is also extended to unknown error distributions by Li (2002) using the repeated-measurement framework discussed in Appendix D.

F.2 Updated Calibration

After estimating the diffusion parameters, we update the calibration to take these values into account. The updated parameters are listed in Table 9, along with the baseline for comparison. We do so in two scenarios: $\sigma_\tau = 0.30$ and $\sigma_\tau = 1$.

F.3 Quantitative Exercise

We now study the gains from the same exercise as in the text. The main results are in Table 10. The first column fixes the wage at its baseline level, isolating the impact of the changing ability distribution. The second column allows the wage to adjust, adding in the additional general equilibrium effect on prices.

Overall, by biasing our parameter β toward zero, our results are a lower bound on the gains at-scale. Figure 12 shows that the same results on the many-to-one relationship between the ATE and at-scale gains holds.

Table 9: Updated Parameter Values

Model Parameter	Description	Parameter (Baseline)	Parameter ($\sigma_\tau = 0.3$)	Parameter ($\sigma_\tau = 1$)
Exogenously varied:				
σ_τ	St. dev. of distortions	0	0.3	1.0
Group 1				
β	Intensity of diffusion	0.538	0.629	0.883
ρ	Persistence of ability	0.595	0.371	0.494
θ	Match technology "quality"	-0.417	-0.326	-0.179
Group 2				
δ	Death rate of firms	0.016	0.016	0.016
σ_0	St. dev. of new entrant ability distribution	0.961	0.961	0.961
ν	Firm bargaining weight	0.5	0.5	0.5
Group 3				
σ	St. dev. of exogenous ability shock distribution	0.75	0.74	0.66
c	Growth factor in ability evolution	-1.92	-2.23	-2.41
ω	Consumption utility weight	0.53	0.54	0.62
α	Ability elasticity in supplier search	0.36	0.36	0.36
η	Ability elasticity in supplier search	0.05	0.05	0.05

Table notes: Group 1 is jointly chosen from the experimental data. Parameters in Group 2 are calibrated to jointly match moments. Group 3 are also set to match baseline data moments, but match 1-1 with target moments. Both are set to match the same set of moments discussed in the main text (see Table 2 for details).

Table 10: Equilibrium Moments

	$\sigma_\tau = 0.3$		$\sigma_\tau = 1$	
	(1)	(2)	(3)	(4)
	Fixed Wage	At-Scale	Fixed Wage	At-Scale
Income	1.08	1.12	1.20	1.38
Ability	1.08	1.14	1.20	1.42
Aggregate Labor Supply	0.92	0.99	0.89	1.00
Wage	1.00	1.14	1.00	1.39

Table notes: All are measured relative to the baseline equilibrium at the give value of σ_τ. Each column reports the new stationary equilibrium after shocking the matching technology, where the first (columns 1 and 3) holds the wage fixed at its baseline level and the second allows it to adjust.

Figure 12: Range of Aggregate Gains for each ATE

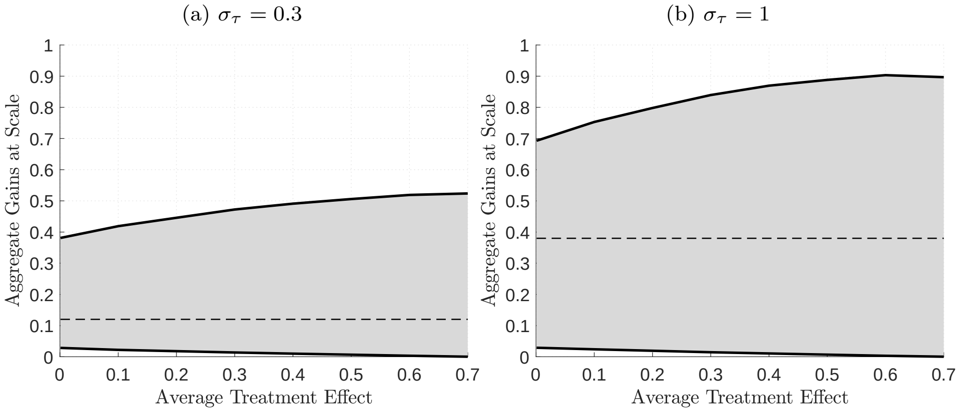

Figure notes: Shaded area is all possible aggregate gains for $(\beta, \rho) \in [0.20, 0.95] \times [0.20, 0.95]$, where θ is chosen to match the ATE on the horizontal axis. The baseline estimate is given by the dashed line.

G Implementing the Same Procedure in the RCT of Cai and Szeidl (2018)

G.1 RCT Details

In an innovative recent paper, Cai and Szeidl (2018) (CS hereafter) conduct an RCT among 2,820 Chinese firms. Treated firms (1,500 of 2,820) are randomly placed into a group of approximately 10 other firms, then compared to a control group with no prearranged meetings.

While similar in style to the RCT used in the main body of the paper, there are a number of differences. First, these are group meetings instead of individual meetings. Second, the meetings are substantially more intense: firms are expected to meet monthly for a year.[35] Third, these firms are larger. At baseline the average firm has 35 workers. Finally, they cross-randomize information about new financial products to track the diffusion of information directly.

Firms are surveyed 3 times: before the treatment, 1 year post-treatment (i.e., at the end of the treatment period), and 2 years post-treatment. We refer to these as $t = 0, 1, 2$. We estimate our procedure off the $t = 1$ data and later ask whether we can match the effect at $t = 2$.

G.2 Overview of Empirical Results

While we will not do justice to the full set of results provided in CS, we provide a brief overview here to relate them back to our baseline RCT in the main text and provide context for modeling decisions. Like our baseline results, CS find important effects on firm performance. In a regression of the form

$$y_{it} = \lambda_0 + \overbrace{\lambda_1 \cdot \mathbb{1}_{t=1} + \lambda_2 \cdot \mathbb{1}_{t=2}}^{\equiv \text{ time effects}} + \overbrace{\lambda_3 \cdot (\mathbb{1}_{t=1} \times T_{it}) + \lambda_4 \cdot (\mathbb{1}_{t=2} \times T_{it})}^{\equiv \text{ per-period ATE}} + FirmFE_i + \varepsilon_{it}$$

they find positive and statistically significant effects at $t = 1$ (i.e., $\lambda_3 > 0$) for sales, profit, employment, and productivity. These meetings do indeed seem to leave firms better off. Unlike our baseline RCT, however, these effects persist at $t = 2$, a full year after the treatment concludes (i.e. $\lambda_4 > 0$).[36]

[35] Take-up is still high, with average attendance of 87 percent.

[36] Productivity here is measured as firm-level TFP from estimating a revenue production function among control firms. This is the only outcome of those listed in which the point estimate at $t = 2$ is statistically insignificant, but it

A number of additional results provide context for these firm performance effects. First, they find persistent changes in management practices.[37] Thus, direct components of firm-level productivity increase. Second, they cross-randomize information on a new financial product to test whether information is indeed flowing between firms within the group, and find that it does. We (and CS) take this as a direct measure that information is flowing between firms in the group.

Finally, CS provide one measure of group-level heterogeneity, asking whether treated firms who have larger average group members enjoy a larger treatment effect. Denoting $\overline{n}_i^m$ as the average firm size of firm i's matches, they run

$$y_{it} = \lambda_0 + \lambda_1 \log(\overline{n}_i^m) + \varepsilon_{it} \quad \text{for } i \text{ in the treatment group.} \tag{G.1}$$

They find statistically significant increases in sales and profit. CS use (G.1) as an "internal consistency check," in the sense that most reasonable theories would predict $\lambda_1 > 0$. As is hopefully clear at this point, this type of regression has an additional role: it is a critical test for extrapolating these results to at-scale implications. We will exploit this regression below in our estimation.

G.3 Model

The empirical setting and results motivate our model structure. Because these are larger firms, we we study this RCT in the context of a more classic Hopenhayn (1992) style model. Given the results on productivity and management practices, we allow diffusion of firm productivity directly, instead of the ability to seek out suppliers. This assumption is more in line with the existing growth literature (Lucas, 2009; Perla and Tonetti, 2014, and many others). Finally, we construct our learning process to conform to the available empirical results, which we discuss more below.

Basics: Production and Households The model period is one year (the length of the exogenous matches in the RCT). There is a measure one of firms, each of which produces according to the production function $y_t = z_t^{1-\alpha-\eta} n_t^{\alpha} k_t^{\eta}$, where z_t is firm-level productivity, n_t is labor, and k_t is capital. Capital and labor are traded on a competitive spot markets with prices r_t and w_t. Each period δ firms exit and are

is still positive. We read less into this result, as it is likely the most noisily estimated.

[37] These include an overall management index, along with components on the evaluation and communication of employee performance, the setting of targets, process documentation, and delegation of power.

replaced by δ new firms, who draw initial productivity from $z \sim G(z)$.[38] As in the main text, firms draw idiosyncratic shocks $\varepsilon \sim F$ and imitation shocks $\hat{z} \sim \widehat{M}$.

A representative household with flow utility $u(C_t)$ provides labor and capital, and owns all firms. Its utility is given by:

$$\begin{aligned} &\max_{\{C_t, K_{t+1}\} \geq 0} \quad \sum_{t=0}^{\infty} (1-\delta)^t u(C_t) \\ &s.t. \quad C_t + K_{t+1} - (1-\lambda)K_t = w_t + r_t K_t + \Pi_t \\ &\qquad\quad K_0 \text{ given} \end{aligned}$$

where Π_t is the aggregate profits from all operating firms and λ is depreciation rate of capital. We assume the household discounts at the firm exit rate for simplicity.

The aggregate state of the economy is the distribution of firm-level productivity, $M(z)$, and the aggregate capital stock, K.

Learning and Diffusion Given the high cost-effectiveness of the program, CS posit a number of possibilities for why firms did not self-organize these meetings. These include search costs, trust barriers and lack of familiarity with other managers, and a free-rider problem in which managers expect others to pay the cost of organization. Motivated by these "missing" meetings, we set up a source distribution in which a firm receives no meetings with probability $1-\theta$. With probability θ, the firm joins a group of exogenous size N. These N firms are uniformly random draws from the equilibrium productivity distribution M. We denote $\hat{z}_1, \ldots, \hat{z}_N$ to be the productivities of the N matched firms.

We next define a match $\hat{z}$ in this context, which in general can take any function over the characteristics of these N firms. Here we are constrained by the results reported in CS, who report how the treatment effect varies by only the average size of the N firms. As such, we assume that $\hat{z} = \sum_{i=1}^{N} \hat{z}_i / N$ so that we can use their provided moment. If a firm does not receive a match, it gets $\hat{z} = 0$.[39] We can therefore write

[38] In the more classic Hopenhayn (1992) or Melitz (2003) sense, the model can be easily extended to include endogenous firm entry/exit by introducing fixed costs, as opposed to our assumed exogenous entry/exit margin. Adding this additional margin does not change the results we focus on here and we therefore exclude it for simplicity. This assumption also does not affect identification. Our procedure relies only on firms in operation at a given point in time, so that the details of past entry are immaterial.

[39] This is of course not a critique of the extremely useful publicly-provided dataset of CS. We note this only to highlight that there is no theoretical benefit to making this assumption, and we do so because it is the only estimating moment of group-level heterogeneity available in their public data.

the source distribution as

$$\widehat{M}(\hat{z}) = 1 - \theta + \theta Q(\hat{z}),$$

where $Q(\cdot)$ is the N-draw sampling distribution of the mean derived from the equilibrium productivity distribution M.

Finally, we note that within treated firms, there is no differential effect between firms with $\hat{z} > z$ or $\hat{z} < z$. Therefore, we remove the max operator and assume that the law of motion takes the form

$$z_{t+1} = e^{c+\varepsilon_t} z_t^{\rho} \left(1 + \frac{\hat{z}_t}{z_t}\right)^{\beta}. \tag{G.2}$$

Equation (G.2) shows that if the firm does not interact that period ($Pr = 1-\theta$), it gains nothing from its $\hat{z}_t = 0$ draw. Conditional on meeting ($Pr = \theta$), however, there will be gains from doing so. Those gains are increasing in the average productivity of matched firms.

Taken together, the firm's problem can be written as (with the aggregate state suppressed, as we will focus on a stationary equilibrium):

$$\begin{aligned} v(z) &= \max_{n,k\geq 0} \; z^{1-\alpha-\eta} n^{\alpha} k^{\eta} - wn - rk + (1-\delta) \int_{\varepsilon} \int_{\hat{z}} v(z'(\hat{z}, \varepsilon; z)) \widehat{M}(d\hat{z}, M) dF(\varepsilon) \\ s.t. \quad & z'(\hat{z}, \varepsilon; z) = e^{c+\varepsilon} z^{\rho} \left(1 + \frac{\hat{z}_t}{z_t}\right)^{\beta} \end{aligned}$$

Equilibrium The stationary equilibrium of this economy is an invariant distribution $M^*(z)$, household decision rules C, K', firm decision rules n, k, and value function v such that the household's and firms' value function solves their respective problems with the associated decision rules, markets clear

1. `labor market:` $\int_z n(z) dM^*(z) = 1$
2. `capital market:` $\int_z k(z) dM^*(z) = K$
3. `consumption market:` $\int_z z^{1-\alpha-\eta} n^{\alpha} k^{\eta} dM^*(z) = C + \lambda K$

and the relevant aggregates are consistent

1. `profit:` $\int_z \pi(z) dM^*(z) = \Pi$

2. `law of motion for ability:`

$$\begin{aligned} M' &:= \Lambda(M(z')) \\ &= \delta G(z') + (1-\delta)\int_0^\infty \int_0^\infty F(\log(z') - \rho\log(z) + \beta\log(1+\hat{z}/z) - c)d\widehat{M}(\hat{z})dM(z) \end{aligned}$$

with $M^*(z) = \Lambda(M^*(z))$.

G.4 Estimating Diffusion Parameters

We begin by using our two-step procedure to estimate the key diffusion parameters (β, ρ, θ).

The production function here generates the usual Cobb-Douglas result that inputs labor and capital, sales, and profit are all proportional to z. Therefore, we are free to use any of these moments for identification (see also Appendix Section D.2 for a further generalization).

Averages are also proportional to z. If a firm matches with N firms of size $\hat{n}_1, \ldots, \hat{n}_N$, then we can infer the average productivity as

$$\hat{n} := \frac{\sum_{j=1}^N \hat{n}_j}{N} = \left(\frac{\alpha}{w}\right)^{\frac{1-\eta}{1-\alpha-\eta}} \left(\frac{\eta}{r}\right)^{\frac{\eta}{1-\alpha-\eta}} \frac{\sum_{j=1}^N \hat{z}_j}{N} \propto \hat{z}. \tag{G.3}$$

Given these results, we focus on firm size here for our estimation. As discussed above, it is the only moment available in the public CS data for the first step of our estimation procedure. Rather than introduce a second dependent variable, we use firm size in the second step as well.

Following our procedure in the text, we will estimate the diffusion parameters using only $t = 0, 1$ data then later check whether the time series matches $t = 2$ outcomes. Thus, the first step of our procedure is to estimate

$$\log(n_i') = c + \rho\log(n_i) + \beta\log\left(1 + \frac{\hat{n}_i}{n_i}\right) + \varepsilon \quad \text{for all } i \text{ in treatment} \tag{G.4}$$

where $\hat{n}_i$ is the average size of matched firms as defined in (G.3). Given $(\hat{\rho}, \hat{\beta})$ we then estimate θ with the same procedure as the main text.

$$\min_{\theta} \quad abs\left(\frac{\mathbb{E}[n_T']}{\mathbb{E}[n_C']} - \frac{\int\int \pi^\rho \left(1+\hat{n}/n\right)^\beta d\widehat{H}_T(\hat{n})dH_T(n)}{\int\int n^\rho \left(1+\hat{n}/n\right)^\beta d\widehat{M}(\hat{n};\theta)dH_C(n)}\right) \tag{G.5}$$

where H_C, H_T are the empirical baseline distributions of control and treatment firms, $\widehat{M}$ is the source distribution for control firms, and $\widehat{H}_T$ is the source distribution for treated firms (given exogenously by the empirical implementation). We measure the empirical ratio in (G.5) as the average treatment effect

$$\log(n_i') = \lambda_0 + \lambda_1 T_i + Controls_i + \nu_i \tag{G.6}$$

where $T_i = 1$ if i is in the treatment group. Table 11 provides the regression estimates for (G.4) and (G.6).

Table 11: Identification Moments

	(1)	(2)
β	0.276	
	(0.053)***	
ρ	0.955	
	(0.028)***	
Treatment		0.076
		(0.044)*
R^2	0.764	0.395
Baseline Control Avg	–	2.694

Table notes: Standard errors are in parentheses. Statistical significance at 0.10, 0.05, and 0.01 is denoted by *, **, and, ***.

Column (1) show that the average firm size is indeed a good predictor of treatment effect magnitude. The ATE in column (2) then implies $\theta = 0.199$. That is, the model infers that it is quite unlikely that such a CS-style group would be created without the intervention.

G.5 Remaining Calibration

We calibrate the remaining model to target parameters similar to the main text, using CS data and other relevant moments. We choose 5 parameters that match directly to their moments. These include a 9 percent firm death rate ($\delta = 0.09$) to match exit rates in China (e.g. Lu, 2021). The standard deviation of new entrant ability matches the standard deviation of log profit for firms that have been open for less than 1 year, which implies $\sigma_0 = 1.23$. The depreciation rate is set to a standard value of $\lambda = 0.06$. Finally, we have the Cobb-Douglas exponents on capital η and labor α. We set $\eta = 0.20$ to match the median firm's baseline capital-output ratio $(rk)/y = 0.20$ then set α to match the median firm's profit-sales ratio, $\pi/y = 0.12$. This implies $\alpha = 0.68$. Finally, we set the standard deviation of the exogenous ability shock to $\sigma = 0.45$ and the productivity drift $c = -1.11$. These two parameters jointly match the standard deviation of log profit in the economy and the ratio of average profit of all firms relative to those with less than 1 year of operation. Parameters and moments are in Table 12.

Table 12: Targets and Parameter Choices

Model Parameter	Description	Parameter Value	Target Moment	Source	Target Value	Model Value
Group 1	*From RCT*					
β	Intensity of diffusion	0.276	Estimated parameter from regression	RCT results	0.276	0.276
ρ	Persistence of ability	0.955	Estimated parameter from regression	RCT results	0.955	0.955
θ	Match technology "quality"	0.199	Treatment effect at $t = 1$	RCT results	0.076	0.076
Group 2	*Matched one-to-one with parameter*					
δ	Death rate of firms	0.09	Average exit rate in China	Literature (Lu, 2021)	34	34
σ_0	St. dev. of new entrant ability distribution	1.23	Variance of log profit among new entrants	CS Baseline	1.23	1.23
α	Cobb-Douglas share, n	0.68	Median firm π/y	CS Baseline	0.12	0.12
η	Cobb-Douglas share, k	0.20	Median firm $(rk)/y$	CS Baseline	0.20	0.20
λ	Depreciation rate	0.06	Literature	–	–	–
Group 3	*Jointly targeted*					
σ	St. dev. of exogenous ability shock distribution	0.45	Standard deviation of log profit in all firms	CS Baseline	1.34	1.34
c	Growth factor in ability evolution	-1.11	Ratio of average profit of all firms to new entrants	CS Baseline	1.12	1.12

Table notes: Group 1 is jointly chosen from the experimental data. Group 2 are also set to match baseline data moments, but match 1-1 with target moments. Parameters in Group 3 are calibrated to jointly match moments.

G.6 Treatment Effect Persistence

In Section **??**, we showed that the treatment effect persistence is increasing in ρ and decreasing in β. In our baseline Kenyan RCT we estimate $(\beta, \rho) = (0.538, 0.595)$. Here, we estimate $(0.276, 0.955)$. Our estimates of both β and ρ predict more persistent treatment effects than our baseline model.

We compare this to the empirics in Figure 13. We plot 3 time paths: the empirical treatment effect (which is available for 2 years post-treatment), the estimated model effect, and the estimated model effect when we assume our Kenyan RCT values for β and ρ. We extend the latter two series for 5 years to trace the dynamics over a longer horizon.

Figure 13: Dynamics of Average Treatment Effect (Firm Size)

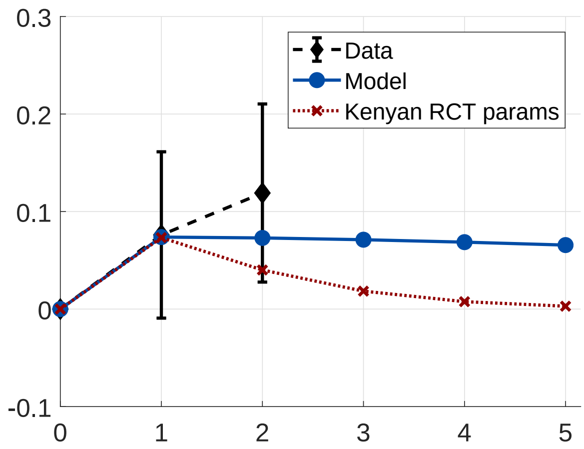

Figure notes: Dashed line is the data with 95 percent confidence interval, for which data is available at $t = 0, 1, 2$. Solid line is the estimated model $(\beta, \rho, \theta) = (0.276, 0.955, 0.199)$. For comparison, we include the results at our Kenyan RCT values of β and ρ, re-estimating θ to hit the $t = 1$ ATE, which implies $(\beta, \rho, \theta) = (0.538, 0.595, 0.496)$. We extend the model-derived RCT for 5 post-treatment years to study fade-out.

The model is consistent with the persistent gains.[40] Even 5 years post-treatment, the model predicts that 89 percent of the initial benefits remain. In comparison, if we replace β and ρ by our estimated values in Kenya, the fade-out is more substantial and nearly complete by $t = 4$. The results highlight the importance of estimating these parameters in governing the time series of the treatment effect.

[40]The slight increase observed in the treatment effect from $t = 1$ to $t = 2$ is statistically insignificant by any reasonable cutoff.

G.7 Quantitative Gains at Scale

We conduct the same exercises as the main text. To measure the aggregate implications, we permanently shock the matching technology to increase average match quality, in line with the extensive margin focus of the RCT results. We do so by shocking the parameter θ so that it is 25 percent closer to its limit of $\theta = 1$. We study the new stationary equilibrium and compare it to the baseline equilibrium. Aggregate moments are reported in Table 13. We present two steady states, both of which operate under the new matching function. The first (in column 2) fixes the wage at its baseline level. The second (in column 3) allows the wage to adjust as well.

Table 13: Equilibrium Moments

	(1) Baseline	(2) Fixed Wage	(3) At-Scale
Income	1.00	1.13	1.05
Ability	1.00	1.39	1.39
Aggregate Labor Supply	1.00	1.00	1.00
Wage	1.00	1.00	1.05

Table notes: Column (1) is the initial equilibrium, normalized to one. (2) and (3) report the new stationary equilibrium after shocking the matching technology, where (2) holds the wage fixed at its baseline level and (3) allows it to adjust.

The direct effect of making it easier to learn from high ability agents is that average ability rises by 39 percent and income by 13 percent. Some of that is eliminated by the higher general equilibrium wage, with the net effect of a 5 percent increase in total household income.

We next ask the importance of measuring treatment effect heterogeneity, as in the main text. We vary $(\beta, \rho) \in [0.20, 0.95] \times [0.20, 0.95]$. For each, we continually update θ to match a given average treatment effect. We vary the ATE from 0 to 18 percent, which traces out the range of possible aggregate outcomes by ATE.[41] Those results are in Figure 14. Similar results emerge as in the main text – the set of possible aggregate outcomes for a given treatment effect can be large.

[41] The set of feasible ATEs for the range of (β, ρ) we consider is smaller here than in the main text, as our diffusion process requires $\theta \in [0, 1]$. See the discussion surrounding Proposition 2 for more details on this constraint.

Figure 14: Range of Aggregate Gains for each ATE (Firm Size)

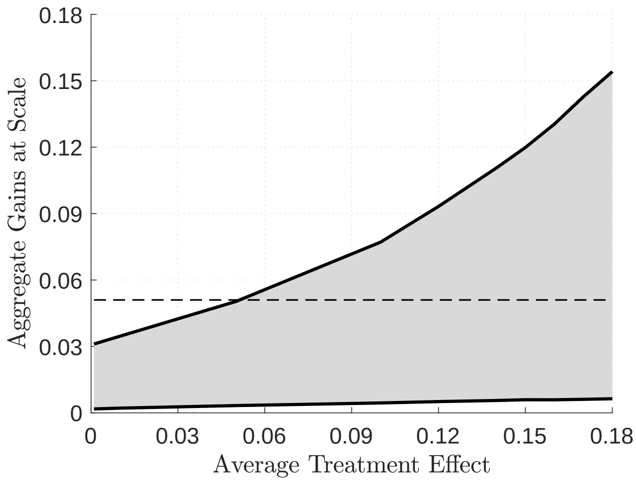

Figure notes: Shaded area is all possible aggregate gains for $(\beta, \rho) \in [0.20, 0.95] \times [0.20, 0.95]$, where θ is chosen to match the ATE on the horizontal axis. The baseline estimate at $(\beta, \rho) = (0.276, 0.955)$ is given by the dashed line.

H Proofs

H.1 Proof of Proposition 2

The bounds used in Proposition 2 are given by

$$\begin{aligned}\Gamma^{min} &= \inf_{\theta} \frac{\int\int \pi^{\rho} \max\{1, \hat{\pi}/\pi\}^{\beta} d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^{\rho} \max\{1, \hat{\pi}/\pi\}^{\beta} d\widehat{M}_{\pi}(\hat{\pi}; \pi, \theta) dH_{C,\pi}(\pi)} \\ \Gamma^{max} &= \sup_{\theta} \frac{\int\int \pi^{\rho} \max\{1, \hat{\pi}/\pi\}^{\beta} d\widehat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\int\int \pi^{\rho} \max\{1, \hat{\pi}/\pi\}^{\beta} d\widehat{M}_{\pi}(\hat{\pi}; \pi, \theta) dH_{C,\pi}(\pi)}.\end{aligned}$$

We proceed with the proof under those bounds.

Proof. We can write $\pi' = g(z, \hat{z}, \varepsilon) = Ae^{c+\varepsilon} z^{\rho} \max\{1, \hat{z}/z\}$ by Assumptions 1 and 2 for some constant A. Since g is continuous, for a density $f(z, \hat{z}, \varepsilon)$ we have

$$\mathbb{E}[\pi'] = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} g(z, \hat{z}, \varepsilon) f(z, \hat{z}, \varepsilon) d\hat{z} dz d\varepsilon$$

This follows from what is sometimes referred to as the "law of the unconscious statistician."[42] Inserting the correct joint distributions, and utilizing the fact that the exogenous shocks are orthogonal to z and $\hat{z}$ (Assumption 1) gives us

$$\frac{\mathbb{E}_T[\pi']}{\mathbb{E}_C[\pi']} = \frac{\int\int z^{\rho} \max\{1, \hat{z}/z\}^{\beta} d\widehat{H}_T(\hat{z}) dH_T(z)}{\int\int z^{\rho} \max\{1, \hat{z}/z\}^{\beta} d\widehat{M}(\hat{z}; z, \theta) dH_C(z)}$$

Applying the proportionality assumption in Assumption 2 yields

$$\frac{\mathbb{E}_T[\pi']}{\mathbb{E}_C[\pi']} = \frac{\int\int \pi^{\rho} \max\{1, \hat{\pi}/\pi\}^{\beta} d\widehat{H}_{T,\pi}(\hat{z}) dH_{T,\pi}(z)}{\int\int \pi^{\rho} \max\{1, \hat{\pi}/\pi\}^{\beta} d\widehat{M}_{\pi}(\hat{\pi}; \pi, \theta) dH_{C,\pi}(z)}.$$

Given β and ρ, the right hand side is continuous in θ by the continuity of $\widehat{M}$ in Assumption 3. The intermediate value theorem then guarantees existence when $\Gamma \in [\Gamma^{min}, \Gamma^{max}]$. Finally, uniqueness follows from the strict monotonicity of the right hand side in θ, which is guaranteed by the assumed first order stochastic dominance in Assumption 3. ■

[42] This result is trivially applied given our assumptions used in the main text, where the equation follows directly from $\pi' \propto z'$. However, it is a useful result when we relax functional form assumptions in various extensions, so we highlight it here.

H.2 Proof of Proposition 3

Proof. Start from the law of motion defined in Assumption 1. In logs, and applying Assumption 2 ($\pi \propto z$) annd Assumption 4 ($\hat{z} > z$), this is

$$\log(\pi_i') = \widetilde{c} + \rho \log(\pi_i) + \beta \log\left(\max\left\{1, \frac{\hat{\pi}_i}{\pi_i}\right\}\right) + \varepsilon_i, \tag{H.1}$$

where π_i and $\hat{\pi}_i$ are baseline profit for individual i in the treatment group and her match, while $\widetilde{c}$ is a constant equal to the structural parameter c if $\pi = z$. Since matches are observable within the treatment, (H.1) is a simple panel regression with coefficients β and ρ. Since matches are randomized, the parameters are identifiable. Thus, that $\hat{\beta}^{OLS}$ and $\hat{\rho}^{OLS}$ are equal to their structural counterparts follows directly from the within-treatment exclusion restriction of Assumption 4. ■

H.3 Proof of Proposition 1

We begin by detailing the underlying arithmetic of the model, then use those results to prove Proposition 1 at the end of this section.

Solving the Bargaining Problem as a Function of Supplier Marginal Cost m Solving the optimal input choices for the firm implies profit is

$$\pi^f(c, w) = \left(\frac{\alpha}{p_x}\right)^{\frac{\alpha}{1-\eta-\alpha}} \left(\frac{\eta}{w}\right)^{\frac{\eta}{1-\eta-\alpha}} (1 - \alpha - \eta). \tag{H.2}$$

A supplier has profit function $\pi^s = (p_x - m)x$ where m is its given marginal cost. Given that they take as given the firm's decision on inputs, this implies we can write this profit as

$$\pi^s(p_x, m) = \left(\frac{\alpha}{p_x}\right)^{\frac{1-\eta}{1-\eta-\alpha}} \left(\frac{\eta}{w}\right)^{\frac{\eta}{1-\eta-\alpha}} (p_x - m).$$

Re-writing the bargaining problem $\pi(p_x)^\nu \pi^s(p_x, m)^{1-\nu}$ taking these derivations into account yields

$$\max_{p_x} \quad \left(\frac{\eta}{w}\right)^{\frac{\eta}{1-\eta-\alpha}} (1 - \alpha - \eta)^\nu \alpha^{\frac{1-\eta-\nu(1-\eta-\alpha)}{1-\eta-\alpha}} c^{\frac{\nu(1-\eta-\alpha)+\eta-1}{1-\eta-\alpha}} (p_x - m)^{1-\nu}$$

Log differentiating gives the solution

$$p_x = \left(\frac{1-\eta-\nu(1-\eta-\alpha)}{\alpha}\right) m \tag{H.3}$$

Replacing p_x in the firm's profit function (H.2) with the value from (H.3) yields

$$\pi = \left(\frac{\alpha}{1-\eta-\nu(1-\eta-\alpha)}\right)^{\frac{\alpha}{1-\eta-\alpha}} (1-\eta-\alpha)\alpha^{\frac{\alpha}{1-\eta-\alpha}} \left(\frac{\eta}{w}\right)^{\frac{\eta}{1-\eta-\alpha}} m^{\frac{-\alpha}{1-\eta-\alpha}} \tag{H.4}$$

Optimal Search Intensity Now that we have profit as a function of the supplier's marginal cost m in (H.4), we need to solve for search intensity s. Recall that $m = exp(-s)z^{\frac{\alpha+\eta-1}{\alpha}}$. Plugging this into (H.4),

$$\pi^f(m) = \left(\frac{\alpha}{1-\eta-\nu(1-\eta-\alpha)}\right)^{\frac{\alpha}{1-\eta-\alpha}} (1-\eta-\alpha)\alpha^{\frac{\alpha}{1-\eta-\alpha}} \left(\frac{\eta}{w}\right)^{\frac{\eta}{1-\eta-\alpha}} exp(-s)^{\frac{-\alpha}{1-\eta-\alpha}} z$$

The relevant part of the firm's utility function for this problem is the static utility flow $\omega \log(\pi) + (1-\omega)\log(1-s)$. Plugging in π and solving for s yields

$$s = 1 - \left(\frac{1-\omega}{\omega}\right)\left(\frac{1-\eta-\alpha}{\alpha}\right) \tag{H.5}$$

Occupational Choice Given the decision rules derived above, the model therefore has a cutoff rule for occupational choice. To see this, note that because the continuation values between workers and entrepreneurs are identical, we can focus on the flow utility payoff. After a bit of algebra, these are

$$u^f(w) = \omega\log(C_1) + (1-\omega)\log(1-s) - \frac{\eta\omega}{1-\eta-\alpha}\log(w) + \omega\log(z)$$

with constants

$$\begin{aligned}
C_1 &= \left(\frac{\alpha}{1-\eta-\nu(1-\eta-\alpha)}\right)^{\frac{\alpha}{1-\eta-\alpha}} (1-\eta-\alpha)\alpha^{\frac{\alpha}{1-\eta-\alpha}}\eta^{\frac{\eta}{1-\eta-\alpha}} exp\left(\left(\frac{1-\omega}{\omega}\right)\left(\frac{1-\eta-\alpha}{\alpha}\right)-1\right)^{\frac{-\alpha}{1-\eta-\alpha}} \\
s &= 1 - \left(\frac{1-\omega}{\omega}\right)\left(\frac{1-\eta-\alpha}{\alpha}\right)
\end{aligned}$$

More simply for workers, flow utility is $u^w(w) = \omega\log(w)$. Firm operation is preferred when $u^f(w) \geq u^w(w)$, which implies a cut-off $\underline{z}$

$$\underline{z} = w^{\frac{1-\alpha}{1-\eta-\alpha}} \exp\left(-\log(C_1) - \frac{1-\omega}{\omega}\log\left[\left(\frac{1-\omega}{\omega}\right)\left(\frac{1-\eta-\alpha}{\alpha}\right)\right]\right)$$

Thus, the cutoff has the feature that $\log(\underline{z}) \propto \log(w)$, where w is the equilibrium wage.

Proof of Proposition 1 With these results, Proposition 1 follows quickly.

Proof. Plugging (H.5) into the profit function yields

$$\begin{aligned}\pi &= \left(\frac{\alpha}{1-\eta-\nu(1-\eta-\alpha)}\right)^{\frac{\alpha}{1-\eta-\alpha}}(1-\eta-\alpha)\alpha^{\frac{\alpha}{1-\eta-\alpha}}\left(\frac{\eta}{w}\right)^{\frac{\eta}{1-\eta-\alpha}} \\ &\quad\times\ exp\left(\left(\frac{1-\omega}{\omega}\right)\left(\frac{1-\eta-\alpha}{\alpha}\right)-1\right)^{\frac{-\alpha}{1-\eta-\alpha}} z.\end{aligned}$$

Thus, we have $\pi = A(w)z$ in equilibrium, where A depends on parameters and the equilibrium wage w. Replacing $m = exp(-s)z^{\frac{\alpha+\eta-1}{\alpha}}$ in the equilibrium input price function (H.3) gives

$$p_x = \left(\frac{1-\eta-\nu(1-\eta-\alpha)}{\alpha}\right) exp\left(\left(\frac{1-\omega}{\omega}\right)\left(\frac{1-\eta-\alpha}{\alpha}\right)-1\right) z^{\frac{\alpha+\eta-1}{\alpha}}$$

as required. ∎

H.4 Proof of Proposition 7 (from Social Planner's Problem)

Proof. Since these are static decisions, they solve the simplified static problem

$$\begin{aligned}\max_{c,s,x,n} &\quad \omega\int_0^\infty \log(c(z))dM(z) + (1-\omega)\int_{\underline{z}}^\infty \log(1-s(z))dM(z) \\ s.t. &\quad \int_{\underline{z}}^\infty x(z)^\alpha n(z)^\eta dM(z) - \left(\frac{1-\eta-\nu(1-\eta-\alpha)}{\alpha}\right)\int_{\underline{z}}^\infty exp(-s(z))z^{\frac{\alpha+\eta-1}{\alpha}}x(z)dM(z) = \int_0^\infty c(z) \\ &\quad \int_{\underline{z}}^\infty n(z)dM(z) = \int_0^{\underline{z}} dM(z) \equiv N_s \\ &\quad M(z) \text{ given}\end{aligned}$$

Note that for simplicity here, we have already imposed a cut-off value for z, $\underline{z}$. Much like the *laissez faire* equilibrium, it is straightforward to show that the planner also chooses to set occupations this way.

The first piece to note is that, given the separability of utility in c and s, the planner allocates a constant level of consumption $c(z) = c$. Thus, we just need to determine total resources in the economy to determine consumption. Solving the optimal input choices $x(z)$ and $n(z)$ collapses the simplified static planner's problem to

$$\begin{aligned}
\max_{s(\cdot)\geq 0} \quad & \omega\log(c) + (1-\omega)\int_{\underline{z}}^{\infty}\log(1-s(z))dM(z) \\
s.t. \quad & \left(\frac{\alpha^2}{1-\eta-\nu(1-\eta-\alpha)}\right)^{\frac{\alpha}{1-\alpha}}(1-\alpha)N_s^{\frac{\eta}{1-\alpha}}\left(\int_{\underline{z}}^{\infty} z\exp(s(z))^{\frac{\alpha}{1-\alpha-\eta}}dM(z)\right)^{\frac{1-\alpha-\eta}{1-\alpha}} = c \\
& M(z) \text{ given}
\end{aligned}$$

The first order condition for this problem is

$$\frac{(1-\omega)m(z)}{1-s(z)} = \lambda C\left(\int_{\underline{z}}^{\infty} z\exp(s(z))^{\frac{\alpha}{1-\alpha-\eta}}m(z)dz\right)^{\frac{-\eta}{1-\alpha}}\left(\left(\frac{\alpha}{1-\eta-\alpha}\right)z\exp(s(z))^{\frac{\alpha}{1-\alpha-\eta}}m(z)\right) - \lambda_2$$

where λ is the Lagrange multiplier on the resource constraint, λ_2 is the multiplier on the non-negativity constraint $s \geq 0$, and C is the constant in front of the integral of the resource constraint. For any z_1 and z_2 with interior solutions ($\lambda_2 = 0$), we have

$$\frac{1-s(z_2)}{1-s(z_1)} = \left(\frac{z_1}{z_2}\right)\left(\frac{\exp(s(z_1))}{\exp(s(z_2))}\right)^{\frac{\alpha}{1-\alpha-\eta}} \tag{H.6}$$

Define for ease of notation $q(z_2/z_1, s(z_1)) = \left(\frac{z_1}{z_2}\right)\exp(s(z_1))^{\frac{\alpha}{1-\alpha-\eta}}(1-s(z_1))$ and the transformation $-t = \left(\frac{\alpha}{\alpha+\eta-1}\right)s(z_2) - \left(\frac{\alpha}{\alpha+\eta-1}\right)$. After some algebra, we can rewrite (H.6) as

$$\left(\frac{q\alpha}{\eta+\alpha-1}\right)\exp\left(\frac{\alpha}{\eta+\alpha-1}\right) = t\exp(t)$$

The solution to this problem is given by the principal branch of the Lambert W function,

$$t = W_0\left[\left(\frac{\alpha}{\eta+\alpha-1}\right)\exp\left(\frac{\alpha}{\eta+\alpha-1}\right)q\right].$$

Undoing the transformation and setting $z_1 = \underline{z}$, if the economy is at an interior solution to s for all $z \geq \underline{z}$ (which we verify at our estimated parameters), we can write this as

$$s(z_2) = \left(\frac{1-\alpha-\eta}{\alpha}\right)W_0\left[\left(\frac{\alpha}{\eta+\alpha-1}\right)\exp\left(\frac{\alpha}{\eta+\alpha-1}\right)q(z_2/\underline{z}, s(\underline{z}))\right] + 1,$$

Since q is decreasing in z_2 and $\eta+\alpha < 1$, that $s(\cdot)$ is increasing follows from the fact that that W_0 is an increasing function. Concavity similarly follows from properties of W_0. Since W_0 is increasing and $x\exp(x)$ is convex, its inverse W_0 is concave.

$\blacksquare$

If there is a corner solution, this analysis remains nearly identical, except one would instead be required to solve for $\hat{z} = argmin_z \ \lambda_2(z) = 0$ instead of $s(\underline{z})$. That is, the z at which supplier search intensity becomes positive. While this is not at issue at our estimated parameters, we of course take this into consideration when counterfactually varying parameters to study how the quantitative implications change.

Appendix References

Brooks, Wyatt, Kevin Donovan, and Terence R. Johnson, "Mentors or Teachers? Microenterprise Training in Kenya," *American Economic Journal: Applied Economics*, 2018, *10* (4), 196–221.

Cai, Jing and Adam Szeidl, "Interfirm Relationships and Business Performance," *Quarterly Journal of Economics*, 2018, *133* (3), 1229–1282.

Härdle, Wolfgang, Hua Liang, and Jiti Gao, *Partially Linear Models*, Springer-Verlag Berlin Heidelberg, 2000.

Hopenhayn, Hugo A., "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Journal of Political Economy*, 1992, *60* (5), 1127–1150.

Hsiao, Cheng, "Consistent estimation for some nonlinear errors-in-variables models," *Journal of Econometrics*, 1989, *41*, 159–185.

Imbens, Guido and Karthik Kalyanaraman, "Optimal Bandwidth choice for the Regression Discontinuity Estimator," *Review of Economic Studies*, 2012, *79* (3), 933–959.

Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg, "Learning from Coworkers," *Econometrica*, 2021, *89* (2), 647–676.

Li, Tong, "Robust and consistent estimation of nonlinear errors-in-variables models," *Journal of Econometrics*, 2002, *110*, 1–26.

Lu, Will Jianyu, "Transport, Infrastructure and Growth: Evidence from Chinese Firms," 2021. Working Paper.

Lucas, Robert E., "Ideas and Growth," *Economica*, 2009, *76* (301), 1–19.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Perla, Jesse and Christopher Tonetti, "Equilibrium Imitation and Growth," *Journal of Political Economy*, 2014, *122* (1), 52–76.

Schennach, Susanne M., "Estimation of Nonlinear Models with Measurement Error," *Econometrica*, 2004, *72* (1), 33–75.

—, "Mismeasured and unobserved variables," in Steven N. Durlauf, Lars Peter Hansen, James J. Heckman, and Rosa L. Matzkin, eds., *Handbook of Econometrics*, 2020, pp. 487–565.

Yatchew, A., "An elementary estimator of the partial linear model," *Economic Letters*, 1997, *57* (2), 135–143.