From Micro to Macro in an Equilibrium Diffusion Model*

Wyatt Brooks
Arizona State University

Kevin Donovan
Yale University

Terence R. Johnson
University of Virginia

March 2024

Abstract

We quantify the benefits of better firm-to-firm matching in an aggregate diffusion model where individuals reap profitable knowledge from others in the economy. We estimate the model using an empirical evaluation of a small-scale program in Kenya that creates new opportunities for firm managers to interact. Critical to the aggregate gains from the program is the relative importance of meeting a high-knowledge firm compared to the learning that happens within that meeting. We show how moments from the intervention identify these diffusion parameters. Doing so formalizes how other easily-estimated moments besides the average treatment effect provide crucial information about at-scale gains. We lastly provide sufficient conditions under which the same identification procedure holds for a wider class of experiments and aggregate diffusion models that includes much recent work.

JEL Classification Codes: O11, O33, O40

Keywords: diffusion, general equilibrium, knowledge

*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, Julieta Caunedo, 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. We are grateful to PEDL (ERG 5817) for financial support.

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 have been 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. The premise of these interventions is that frictions limit interactions with more skilled or knowledgeable firms, which lowers managerial capacity and thus firm growth. Facilitating meetings between firm managers, and thus overcoming these frictions, would therefore be beneficial. Results from several studies support this view, and show that interventions encouraging learning from other managers increase profit, technology adoption, and managerial skills.[1]

Yet the general equilibrium, at-scale implications of these interventions are less explored, and there are reasons to suspect they may differ from what can be extrapolated directly from these treatment effects. First, general equilibrium forces can complicate the link from treatment effects to at-scale outcomes (Buera et al., 2021a). Here, this takes the form of knowledge spillovers. When the knowledge created in these interventions can further permeate the economy through diffusion, the aggregate gains from better matching will depend on changes to the entire distribution of knowledge (Lucas and Moll, 2014; Perla and Tonetti, 2014). Second, the benefits of making it easier to meet high-knowledge agents will depend on a variety of factors. How hard is it to meet these agents without an external intervention? How much can be learned once those meetings do occur? If these "meeting" and "learning" technologies have different contributions to aggregate

[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. (2020) randomly seed information on new technology in Malawi. See also Giorcelli (2019) and Bianchi and Giorcelli (2019) for historical examples of firm-to-firm knowledge transfers and Munshi (2007) and Breza et al. (2019) for reviews.

outcomes, it creates the possibility that the same average treatment effect is consistent with many possible aggregate outcomes.

Our contribution in this paper is two-fold. We formalize the link between experimental moments and elements of the diffusion process in a general equilibrium model of firm-to-firm diffusion. These moments provide a natural way to separate the relative importance of meeting and learning technologies. We then show that doing so matters quantitatively both for understanding aggregate implications and for interpreting reduced-form moments for the ultimate goal of measuring of at-scale policy potential.

We focus our discussion around the results of a randomized controlled trial in Kenya (Brooks et al., 2018) which we discuss in Section 2, though we show many of the insights can be formally generalized to a class of models and interventions that cover a larger body of recent work. In this RCT, we randomly paired high- and low-profit female business owners in Nairobi, Kenya. There are three main findings. First, 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. Second, the treatment effect is largest among those who are randomly paired with higher-profit matches. Third, the underlying mechanisms are primarily on the cost side: treated owners were more likely to switch suppliers and their costs declined significantly.

Our RCT, like others cited above, introduces a one-time shock in partial equilibrium. By that, we mean it focuses on the direct effect of those engaged in the match, and does so at a small enough scale that spillovers to the control group are negligible. Our interest is understanding how that relates to our ultimate policy goal: the introduction not of a one-time change, but a systematic, permanent policy that makes it easier for everyone to learn. Doing so requires us to confront two issues: (1) the difference between a short- and long-run change and (2) the introduction of general equilibrium forces that arise only at scale.

To confront these issues, we build a general equilibrium model of knowledge diffusion and use it as a laboratory to study these questions. We break the paper into two parts. This first part, detailed in Section 3, focuses exclusively on how to map moments from the one-time RCT shock to model parameters. We do so under the minimal amount of necessary structure to clarify how our assumptions let us link model parameters to empirical moments. We devise a simple estimation strategy then provide sufficient conditions on both the general equilibrium model and partial equilibrium experiment that guarantee the procedure identifies the relevant parameters. We view this as a useful contribution because there are many models of diffusion and a growing set of RCTs that link social meeting to knowledge transfer, but differ in their exact implementation.[2] The results here clarify what assumptions are directly required for parameter estimation without obfuscating with additional detail.

The identification procedure boils down to separating the processes governing how people learn from a given meeting from how likely they are to enter into that meeting in the first place. For the former, we use the fact that treatment intensity (the quality or profitability of a match) is random. Thus, the more correlated mentor profit is with the individual-level treatment effect, the easier it must be to learn. In practice, this requires estimating an adjusted covariance moment between baseline mentor profit and *ex post* mentee profit. For the latter, we then adjust the distribution of individual-level treatment effects until the average is correct. Thus, the first step uses heterogeneity within the treatment, while the second uses the average treatment effect. The generality of this procedure follows from the same orthogonality conditions used to measure causal treatment effects. We only have to specify certain parts of the model for estimation, leaving the remaining model

[2]For example, while our experiment and Lafortune et al. (2018) use one-to-one meetings, Cai and Szeidl (2018) introduce group meetings of firm managers and Atkin et al. (2017) vary buyer-supplier links.

structure to be adjusted to the specifics of the economic setting considered by the researcher.

The second part of the paper, Sections 4 and 5, quantifies the significance of this link. Specifically, we ask how average income changes when there is a permanent change that makes it easier to meet high knowledge agents. We view this as the natural policy progression taking partial equilibrium effects measured in an RCT to "at scale" policy by institutionalizing the policy change. This allows the entire economy easier access to high-knowledge matches, rather than a small subset where knowledge spillovers are presumed to be negligible. We build a model that allows for general equilibrium gains created knowledge diffusion — a potentially important force absent from the RCT results — and use it as a laboratory to investigate the long-run and aggregate impact of the policy. We implement this by shocking the meeting technology. For a rough sense of magnitude, the shock we induce increases the knowledge in the average meeting by 42 percent between the two steady states in our calibrated model. We refer to this shock as a "policy change," imagining that policy can target this parameter of the economy and facilitate more productive diffusion of business knowledge.

There is no mechanical reason why the model needs to deliver large income changes from this policy change. Varying the estimated parameters with reasonable ranges can generate widely varying results: anywhere from a 0 to 40-fold increase in average income. We show that the model delivers large gains when two complementary forces operate. First, the policy must induce a large change in meeting quality. For a fixed policy shock, that amounts to a baseline economy in which it is hard for agents to meet good matches. Second, agents must learn a lot once they meet a good match. This learning ability allows good meetings to translate into more knowledge, which can be perpetuated in tomorrow's meetings with others. Taken together, large aggregate gains require a baseline economy characterized by both a poor meeting technology and a good learning technology. This complementarity, and how various empirical moments relate to it, is central to most of the quantitative results discussed in this paper.

We estimate the learning and meeting parameters using our approach described above, then offer a relatively standard calibration for the remaining model structure. We find that average income increases by 11 percent in the new steady state. Sixty percent of that effect comes directly from making it easier to meet high-knowledge firms. Because everyone can more easily accumulate knowledge, the equilibrium stock of knowledge increases and thus increases average income. The remaining 40 percent comes from an amplification effect through price changes. More profitable firms demand more workers and therefore pay higher wages. This competitive pressure pushes marginal firms out of business and reduces congestion in learning for everyone else, further expanding the stock of knowledge.

These results comprise our first quantitative contribution in this paper: we show how moments from a set of successful recent experiments inform otherwise hard-to-measure parameters in an aggregate diffusion model, and that doing so matters quantitatively for measuring impact of scaling the experiment. Our second contribution is to link this relationship to a more micro-development question. We ask when, if ever, it makes sense to extrapolate aggregate implications from the impact on the average firm in the experiment. This "average treatment effect" or ATE is often used as a measure of success in an experiment and as motivation for scaled policy (e.g. World Bank, 2020). But the way we derive model parameters uses both this and other moments together. To what extent is this extrapolation from the average effect an innocuous one?

We use the model to show that the ATE directly informs the the general equilibrium at-scale gains only when the ATE is small. That is, a small ATE implies small at-scale gains, but a large ATE does not guarantee large at-scale gains. In this case, one needs to use the covariance moment we highlight to clarify among a set of possible outcomes. We show that this parameter can be estimated with a simple linear regression.

The intuition for this result follows from the relationship between meeting and learning. As discussed above, aggregate gains depend on their complementarity. It requires that agents find it difficult to meet good matches but easy to learn from them once they do. But this cannot be true if the ATE is small. In this case, the baseline economy must be characterized by (1) a meeting technology such that everyone can already meet high-quality matches or (2) a learning technology such that no one can learn once they do. The aggregate complementarity implies aggregate gains are small in either scenario. A small ATE therefore implies small gains from scaling the intervention. A large ATE, on the other hand, can be rationalized by a larger set of meeting and learning parameters. These different parameter combinations give way to substantially different aggregate outcomes. Here, the covariance moment we estimate allows us to pin down where in this set of possible outcomes we fall. Thus, this empirical moment plays a particularly useful role in understanding aggregate outcomes exactly when the ATE is largest.

The quantitative magnitudes can be large. While a 1 percent ATE is consistent with aggregate gains between 0.005 and 0.1 percent, a 100 percent ATE is consistent with gains between 10 and 283 percent. At the ATE we estimate in our experimental results, the model allows for the possibility of aggregate gains between 5 and 40 percent. As such, measuring covariance provides useful information about at-scale potential at empirically relevant levels of the treatment effect. Its ease of measurement makes it a particularly straightforward moment for researchers or policymakers to collect as part of a partial equilibrium RCT.

Related Literature This paper join a recent literature that disciplines macro-development models with experiments to simulate scaled policy (e.g. Buera et al., 2021b; Caunedo and Kala, 2022;

Fried and Lagakos, 2022; Fujimoto et al., 2021; Kaboski et al., 2022; Lagakos et al., 2018). We further show how measuring additional moments within the RCT provides relevant information about the gains at scale that are useful for policy decisions.[3] 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 or knowledge 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.[4] Our focus on supplier churn and knowledge depreciation relates to Baqaee et al. (2024) on the growth consequences of supplier churn. Finally, related to our focus on the complementarity between meeting and learning, Van Patten (2020) and Jones (2022) highlight the importance of different knowledge production functions on various macroeconomic outcomes.

2 Empirical Evidence on the Benefits of Meetings

Recently, several microeconomic studies have documented the potential benefits for firm owners or managers who are randomly chosen to interact with highly skilled individuals or participate in supply chains. We describe a variety of other studies in the Appendix — all of which are well-executed and provide important insights into social diffusion of entrepreneurial knowledge — but focus mainly on the randomized controlled trial in Brooks et al. (2018) since it is both typical of this literature and we have access to the relevant data. Space constraints naturally require us to leave out some details, but we refer interested readers to Brooks et al. (2018) for other details that

[3] An alternative would be to run larger and larger clustered experiments to measure equilibrium effects. 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.

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

are less relevant for subsequent estimation and modelling.

2.1 Details of RCT and Data Collection

The experiment took place in Dandora, Kenya, a dense informal settlement on the outskirts of Nairobi with over 150,000 residents occupying approximately four square kilometers. Despite rampant poverty, some businesses thrive, and this motivated the idea to match more profitable entrepreneurs with less profitable ones as an alternative to more traditional business classes. In brief, we survey a random sample of 3,290 firm owners that mirrored the characteristics of the population, assign younger and less profitable entrepreneurs to a mentor treatment group or a control group at random, recruit the most profitable firms with at least five years of experience into the study as mentors, match the mentors to the mentees at random, and then survey all study participants over 17 months to measure changes in business practices and profitability over time.

The treatment and control groups are composed of female business owners who have been in operation for less than 5 years[5]. Characteristics across the control and treatment groups are balanced on the randomization, and details are provided in Appendix A.

The mentors were selected from the set of businesses with owners over 40 years old and at least 5 years of experience. From those established businesses, we recruited business owners with the highest profit until we had a sufficient number for matching to our treatment group. 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 than the population average and the mentors are drawn from the far right tail. Firms were then surveyed over 6 quarters to track the time series of treatment,

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

although our analysis will focus primarily on data from the baseline and first wave.

Anticipating the model and quantitative results, key for our identification is the two layers of randomization. The first is the usual randomization between control and treatment, and the second is the randomization of individual matches within the treatment group. The first layer will allow us to use variation between treatment and control to identify equilibrium parameters of the economy, while the second layer will allow us to identify the rate of social diffusion of entrepreneurial knowledge.

Details of a "Match" Any randomized controlled trial must balance the scope and scale of the intervention with external validity: The mentor-mentee relationship was intended to replicate important features of real world social diffusion of entrepreneurial knowledge and still appeal to participants to ensure take-up. 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." We provided no topics to discuss, preferring that the content was self-directed. We offered mentors only some optional, vague prompts in an initial orientation meeting that could be used to start a conversation (e.g. "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 and many continued meeting. We placed no artificial restrictions on who mentors could provide business advice to, or from whom mentees could ask for help.

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] We designed the details to be

[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 as well as fixed, relatively large buildings from which they conducted business. Moreover, the first meeting took place at the

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).

2.2 Treatment Effects and Underlying Mechanisms

Figure 1 plots the average treatment effect (as a percentage) over time, along with the 95 percent confidence interval. There is a large average treatment effect in the immediate aftermath of the first meeting that fades over the subsequent three quarters. In previous work, we showed that this could largely be explained by dissolution of matches over time: As pairs broke down, the Average Treatment Effect declined, but the Treatment on the Treated among pairs that continued to meet remained consistent with impacts in early periods. For this reason, we interpret the impacts of a mentorship as a short-run benefit that gradually reverts to the long-run average. We test whether our model rationalizes this pattern after laying out the model below, and find that it does.

Figure 1 plays a key role in the subsequent analysis. The magnitude of the wave one effect of mentorship is a 46.9 percent increase in profit in the mentorship group relative to the control group, which we denote throughout the paper as $\text{ATE}^{\text{data}} = 1.469$.

In general, we do not use the data from waves 2 through 5 to estimate key parameters. In principle, these data could be used to refine our estimates of micro-parameters that govern the diffusion and persistence of entrepreneurial knowledge. In practice, however, we only observe so much information about firms over time in later waves, including with whom they meet and share knowledge. Rather than using a structural approach that explicitly models the durability of matches and unobserved meetings and then integrates out the unobserved values, we restrict attention to a small number of periods in which the intervention is the dominant factor in social diffusion of knowledge. Ultimately, this simpler approach can be validated by comparing the fade-out in the

mentor's business. Thus, that the mentor was "good" at running a business would likely have been understood with or without us.

treatment effect using the estimated parameters, which we provide in Section 4.3.

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 focus on further modeling and estimating 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, our model estimation will imply this as well.

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]

3 What do the empirics represent in a diffusion model?

In this section, we define and estimate a model of diffusion. We show that the key parameters of this diffusion model can be estimated independently from broader structure of the model.[8] We view this as useful for two reasons. First, it lets us formalize sufficient conditions under which the link between the RCT moments and model parameters holds. This allows researchers to use our approach without need to follow the exact model we lay out later in Section 4. We view this as particularly useful given the wide variety of interventions that show that entrepreneurial

[7] See Brooks et al. (2018) for details on the procedure. This forms the basis for our assumption of the max operator in the model we use below. We also observe no loans, joint input purchases or bulk discounts, profit sharing, or other mechanisms that suggest alternative theories.

[8] For example, we can estimate the parameters of the diffusion process without need to specify how prices are determined in equilibrium, or whether or not the economy is in a stationary equilibrium.

knowledge diffuses through social contact. Second, doing so clarifies how specific assumptions on model environment interact with the randomization used to derive the empirical estimates. This section discusses the relationship between model parameters and the RCT estimates while Sections 4 and 5 then turns to quantifying the importance of this relationship.

3.1 Required Assumptions on the Economic Environment

We assume that each entrepreneur i has knowledge z_i. This index of knowledge is an input into business operations, and fluctuates over time. There is a purely random component to these fluctuations reflecting entrepreneurs acquiring new skills or market-specific information, but also that some of their knowledge may become obsolete or irrelevant.

Entrepreneurs can also learn from others, mirroring a reality where social learning is a crucial component of human capital accumulation. We call such a situation an *imitation opportunity*. If an entrepreneur i with knowledge z_i meets another entrepreneur with knowledge z_j, the imitation opportunity that period for entrepreneur i is $\hat{z}_i = z_j$.

We assume that the evolution of entrepreneurial knowledge follows a process in which current knowledge, imitation opportunities, and individual shocks drive the dynamics of knowledge as summarized in Assumption 1.

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

$$z'(z,\varepsilon,\hat{z}) = e^{c+\varepsilon} z^{\rho} \max\left\{1, \frac{\hat{z}}{z}\right\}^{\beta}, \tag{3.1}$$

where the parameter c is a constant, β is diffusion intensity, and ρ is persistence. ε is uncorrelated with z and $\hat{z}$.

This presumes a process in which an agent with knowledge z gets the opportunity to imitate an agent with knowledge $\hat{z}$. The rate of this exchange is governed by β, while ρ controls the rate of own-knowledge depreciation. If $\beta = 0$, imitation opportunities have no effect, and if $\rho = 0$, there is no persistence. On the other hand, if β and ρ equal 1, the process simplifies to $z' = e^{c+\varepsilon} \max\{z, \hat{z}\}$, and all knowledge is transferred from the more to the less knowledgeable entrepreneur. The max operator implies that the flow of entrepreneurial knowledge goes only in one direction, from the more knowledgeable entrepreneur to the less knowledgeable one[9].

The other parameters of this process have a standard AR(1) interpretation. The constant c corresponds to the long run average knowledge that an entrepreneur would acquire in a world without any knowledge diffusion. Finally, the shock ε represents the depreciation of an individual entrepreneur's knowledge through obsolescence or irrelevance when negative, or accumulation of new knowledge through experimentation or happenstance when positive.

Our second assumption addresses the fact that entrepreneurial knowledge is not directly observable. We assume that more profitable entrepreneurs are more knowledgeable, allowing us to infer relative knowledge from relative profits.

The next two assumptions are on equilibrium outcomes of the model. The first relates unobservable knowledge z to observable characteristics.

Assumption 2. *Profits are homogeneous of degree one in knowledge, so* $\pi(z_i) = z_i \pi(1)$.

The critical feature in Assumption 2 is that there exists some monotonic way to index an unobserved and conceptual quantity like knowledge, z, to a measurable and observable quantity like profit, $\pi(z)$. Any multiplicatively separable version of this assumption where $\pi(z) = \phi(z)\pi(1)$ can be simplified to this one merely by reinterpreting $\tilde{z}' = \phi(z)$ and rescaling the latent knowledge vari-

[9] We show later that this is not a critical assumption but is supported by empirical evidence (Brooks et al., 2018; Jarosch et al., 2020).

able.[10] More generally, Appendix C shows that this can be relaxed to non-separable relationships that include production function estimation based on semi-parametric methods.[11] An additional motivation is that many models already exhibit homogeneity or multiplicative separability properties in equilibrium, such as those based on Cobb-Douglas or constant returns to scale production technologies.

Finally, the third assumption parameterizes the source distribution from which imitation opportunities $\hat{z}$ are 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. ***Continuity:*** *For every z and $\hat{z}$, $\widehat{M}$ is continuous in θ*

2. ***First Order Stochastic Dominance:*** *If $\theta_1 < \theta_2$, then for all $\hat{z}$ and z, $\widehat{M}(\hat{z}; z, \theta_2) \leq \widehat{M}(\hat{z}; z, \theta_1)$.*

It is useful to pause on Assumption 3 and understand what it requires. It requires the researcher to assume an equilibrium source distribution $\widehat{M}$ up to a parameter θ. There is nothing in the RCT of Section 2 that helps identify the type of network or matching function by which people interact. Once a functional form for $\widehat{M}$ is specified, our approach offers a solution to estimating an unknown parameter θ of that function.

Yet, Assumption 3 remains quite flexible. It can handle alternatives ranging from random search models to pure assignment models. In the former θ can be thought of as learning from the maximum order statistic of a random sample of other entrepreneurs or the number of draws

[10] We use such a normalization later to transform the equilibrium of our model so that profits are linear in z.

[11] 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-invariables models (see Schennach, 2020, for a thorough review).

per period. In the latter case, θ represents a measure of complementarity between one's own knowledge and her imitation opportunity. Our source distribution also allows for the possibility that agents with different z draw from different distributions. We discuss various alternatives that fall under this assumption in the Appendix.

Finally, we emphasize that Assumptions 2 and 3 are statements only about the characterization of the equilibrium outcomes from the model, not the micro-foundations that generate them. This implies that these 3 assumptions characterize diffusion in a variety of models, including those summarized in Alvarez et al. (2008).[12]

To recap the key assumptions made here, we have made assumptions that govern two aspects of the model. The first is what we refer to as *learning*: the parameters β and ρ govern how much an agent learns conditional on meeting another agent with knowledge $\hat{z}$. The second is about *meeting:* the parameter θ governs the likelihood of meeting someone with high $\hat{z}$. We build both forces into our model for two main reasons. First, because they are conceptually distinct economic forces. Second, when we want to use the model as a laboratory to study policy experiments, the quantitative predictions will depend critically on their values. Because of that, our next step is to develop a strategy to separate them.

3.2 Intuition for Identification Using RCT Results in Section 2

Given these three assumptions, our goal is to show how three key parameters (β, ρ, θ) are related to moments from our RCT. In this sub-section, we build intuition by walking through the steps using the RCT from Section 2. We then formalize the procedure under more general instruments in Section 3.3. Before getting into identification, we first lay out how we view the RCT in light of

[12]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.

these three assumptions.

Defining the RCT as a Shock to Model Parameters Our interpretation of the RCT is that it that it represents a one-time shock to the source distribution. The control group receives imitation opportunities drawn from $\widehat{M}(\hat{z}, z, \theta)$. Our intervention replaces this source distribution with an exogenously defined $\widehat{H}_T(\hat{z})$ for one period. Moreover, because the individual matches are randomized, the mentor knowledge $\hat{z}_i$ is orthogonal to baseline treated firm knowledge z_i for all i in the treatment.

Our procedure makes use of both the treatment-control randomization and the within-treatment randomization. First, we estimate β and ρ using the latter. Second, we estimate θ as a function of those two parameters using the former source of variation.

Estimating β and ρ Our first step is to identify (β, ρ) in equation (3.1). Taking logs of that equation, and utilizing the fact that $\pi \propto z$, we can re-write (3.1) as

$$\log(\pi') = \tilde{c} + \rho \log(\pi) + \beta \log\left(\max\left\{1, \frac{\hat{\pi}}{\pi}\right\}\right) + \varepsilon. \tag{3.2}$$

Estimating parameters in (3.2) entails a variety of potential challenges. First, imitation opportunities are typically unobserved. Second, strategic search for superior entrepreneurs likely creates selection, implying that ε and $\hat{\pi}$ are correlated and exhibit simultaneity bias. Third, even if the previous issues are overcome, naive OLS estimates will be biased due to the autocorrelation of profit over time.

This is where the experimental nature of the data becomes valuable. Denote the set of treated firms by $\mathbf{T}$, and focus only on the quarter in which the first meeting between the mentor and mentee occurred. By virtue of the experimental design, mentors are more profitable than mentees, and we know that $\hat{\pi}_i \geq \pi_i$ for every treated firm. Together, they imply (3.2) can be more simply written as

$$\log(\pi_i') = \tilde{c} + \rho \log(\pi_i) + \beta \log\left(\frac{\hat{\pi}_i}{\pi_i}\right) + \varepsilon_i \quad \text{for } i \in \mathbf{T}. \tag{3.3}$$

We know that the imitation opportunities $\hat{\pi}_i$ are randomly assigned and there is no selection or omitted variables bias with respect to $\hat{\pi}$. Thus, the within-treatment randomization of matches eliminates the potential simultaneity bias. Moreover, by Assumption 1, last quarter's profit π_i is uncorrelated with the shock ε_i. These facts jointly imply that (3.3) provides unbiased estimates of ρ and β when run as a simple linear regression on one wave of data.

Estimating θ The next step is to estimate the shifter in the source distribution, θ. To start, define the ratio of average *ex post* profit for treatment and control firms, which we refer to as $\text{ATE}^{\text{model}}$ (the average treatment effect, in percentage terms, implied by the model). This requires simply integrating the law of motion over all possible baseline profit - mentor profit - exogenous shock combinations,

$$\begin{aligned}
\text{ATE}^{\text{model}} = \frac{\mathbb{E}[\pi_i' | i \in \mathbf{T}]}{\mathbb{E}[\pi_i' | i \in \mathbf{C}]} &= \frac{\mathbb{E}_{\pi,\hat{\pi},\varepsilon}\left[e^{\varepsilon + c} \pi^{\rho} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\beta} \middle| i \in \mathbf{T}\right]}{\mathbb{E}_{\pi,\hat{\pi},\varepsilon}\left[e^{\varepsilon + c} \pi^{\rho} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\beta} \middle| i \in \mathbf{C}\right]} \\
&= \frac{\mathbb{E}_{\pi,\hat{\pi}}\left[\pi^{\rho} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\beta} \middle| i \in \mathbf{T}\right]}{\mathbb{E}_{\pi,\hat{\pi}}\left[\pi^{\rho} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\beta} \middle| i \in \mathbf{C}\right]}
\end{aligned}$$

where the first line uses the proportionality of profits and knowledge, and the second utilizes the fact that the exogenous ε shocks are uncorrelated with knowledge and thus integrate out.

Letting $H_{T,\pi}$ and $H_{C,\pi}$ be the distribution functions of baseline profits in the treatment and control group respectively, we can write the population moment that corresponds to the average treatment effect from the RCT as

$$\mathrm{ATE}^{\mathrm{model}} = \frac{\iint_{\pi,\hat{\pi}} \pi^{\rho} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\beta} d\hat{H}_{T,\pi}(\hat{\pi}) dH_{T,\pi}(\pi)}{\iint_{\pi,\hat{\pi}} \pi^{\rho} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\beta} d\widehat{M}_{\pi}(\hat{\pi};\pi,\theta) dH_{C,\pi}(\pi)}. \tag{3.4}$$

Notice that the right-hand side is entirely observable once β and ρ are estimated (using equation (3.3) above), except for θ. The functional form $\widehat{M}$ is assumed and, combined with the assumptions above, the right hand side is made up entirely of observed profit distributions.

From there, we can estimate θ by choosing θ so that $\mathrm{ATE}^{\mathrm{model}} = \mathrm{ATE}^{\mathrm{data}} = 1.469$. That is, the average treatment effect implied by running the identical RCT in the model generates the empirical average treatment effect. Formally, it follows almost immediately from the first order stochastic dominance assumption in Assumption 3. Economically, higher θ means that the control group can find better matches. Thus, the higher the baseline θ, the smaller average value offered to the treatment group when they are guaranteed a good match from $\widehat{H}_T(\hat{\pi})$.

To understand how θ is identified here, consider fixed parameters $\beta, \rho > 0$ and a given observed average treatment effect from the data, ATE^{data}. Suppose the treatment assigns imitation opportunities from high profit firms. If the average treatment effect in the data is small (say, close to zero), our procedure would match that fact by estimating a high value for θ, which makes the denominator in equation (3.4) large and therefore ATE^{data} close to one. Economically, this means that the draws experienced by firms in the control group are of similar quality to those received in the treatment group, so if treatment imitation opportunities are high, then θ is large. On the other hand, the larger is the observed treatment effect in the data, the lower is the estimated value of θ, since the model rationalizes large treatment effects as low imitation opportunities by firms in the control group.

While estimating β and ρ depends on the somewhat less standard randomization within the treatment group, the ability to remove unobservables here is a function of the more classic control/treatment randomization – it guarantees the exogenous shocks ε play no role. We next offer a formalization of this procedure under more general conditions.

3.3 Formalizing the Procedure Under More General Conditions

The previous results focus specifically on the RCT we ran. But many other experiments have a similar focus but different implementation (e.g. Atkin et al., 2017; Cai and Szeidl, 2018; Lafortune et al., 2018; Fafchamps and Quinn, 2018, among others).[13] Here, we formally define an experiment as a data generating process which provides the relevant characteristics necessary for the intuitive procedure above to hold. We lay this out in Assumption 4.[14]

Assumption 4. *A set of agents with profit $H_\pi(\pi)$ are observed in two consecutive periods. That set is partitioned into two subsets* **C** *and* **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) ***Randomization:*** *An individual's inclusion in* **T** *and* **C** *is orthogonal to unobserved characteristics.*

(b) ***Control firms continue to meet:*** *We cannot observe the imitation opportunity for any agent in* **C***. They are drawn from a distribution $\widehat{M}_\pi(\hat{\pi};\pi,\theta)$.*

[13] 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.

[14] 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 to understand how a particular type of experiment that has been considered in the literature relates to a class of diffusion models outlined in Section 3.1. Assumption 4 is the formalization of the way in which we tie our hands to a particular type of variation.

(c) ***Treated imitation opportunities are observable, non-null learning occurs:*** *We observe the imitation opportunity for agents in* $\mathbf{T}$*. The distribution of those profits are denoted* $\widehat{H}_{T,\pi}(\hat{\pi}) \neq \widehat{M}(\hat{\pi};\pi,\theta)$*. 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) ***The exact imitation opportunities are random:*** *A treated individual's exact imitation opportunity is orthogonal to unobserved characteristics*

Assumption (a) is the usual exclusion restriction. It puts restrictions on how the unobserved characteristics vary between control and treatment. (b) formalizes the intuitive notion that we cannot observe individual-level imitation opportunities for control agents, but that matches continue to happen in the background of the economy. That is, the control group does not stop receiving imitation opportunities because of the intervention.

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

Our experiment in Kenya satisfies these conditions. We randomize into control and treatment to satisfy (a). (b) is assumed. The randomization of the exact matches within the treatment satisfies (c) and (d). As discussed above, these two layers of randomization are central to the results.

Given this, Proposition 1 lays out the same identification procedure discussed above. However, it substitutes Assumption 4 for our specific RCT.

Proposition 1. *In any model equilibrium that satisfies Assumptions 1, 2, and 3, any instrument (or set of instruments) that satisfies Assumption 4 provides consistent estimates of the three parameters*

(β,ρ,θ) via the following three-step estimation procedure.

1. *Collect coefficients (β,ρ) from the regression*

$$\log(\pi_i') = \tilde{c} + \rho\log(\pi_i) + \beta\log\left(\max\{1, \frac{\hat{\pi}_i}{\pi_i}\}\right) + \varepsilon_i \quad \textit{for } i \in \mathbf{T}$$

2. *Using the estimates (β,ρ) from Step 1, check whether the empirical ATE falls within the relevant range, $ATE^{data} \in [\Gamma^{min}, \Gamma^{max}]$, where*

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

   If $ATE^{data} \notin [\Gamma^{min}, \Gamma^{max}]$, there is no parameter θ that can match ATE^{data} given the estimates (β,ρ) from Step 1. If $ATE^{data} \in [\Gamma^{min}, \Gamma^{max}]$ proceed to Step 3.

3. *If $ATE^{data} \in [\Gamma^{min}, \Gamma^{max}]$ there is a unique value of θ that solves*

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

3.4 Discussion of More General Result

Proposition 1 formalizes the link between model parameters and a class of randomized controlled trials that link social interaction to knowledge transfer. But how do those parameters vary with these moments? And what do they represent economically? We discuss those issues here.

First, notice that Step 3 of Proposition 1 shows that there is a function $\theta(\beta,\rho)$ that can rationalize the average treatment effect. That is, there is always a single θ that rationalizes the ATE for any values (β,ρ) within the empirically-relevant bounds. We refer to this function $\theta(\beta,\rho)$ as the *meeting-learning frontier*, because θ governs "meeting" while (β, ρ) govern "learning" conditional on meeting. We start with an empirically-relevant corollary in Corollary 1.

Corollary 1. *If treatment matches $\widehat{H}_\pi(\hat{\pi})$ first order stochastically dominate control matches $\widehat{M}_\pi(\hat{\pi}; \pi, \theta)$ for any π (that is, the intervention offers better matches than agents usually receive) then $\partial\theta/\partial\beta >$ 0 on the set for which θ is defined.*

Corollary 1 says that the meeting-learning frontier restricts to parameter combinations in which the ease of learning and meeting positively co-move. While the proof follows closely from the FOSD assumption, the economic intuition follows from model's ability to trade off learning versus meeting when matching ATE^{data}. Because easier learning (higher β) increases the average value of the treatment in the model, the model infers it must be easy for agents to find high-knowledge matches without the intervention (higher θ). This lowers the value of the intervention: when agents can easily find high-knowledge agents without the intervention, the value of offering them one exogenously is of little value. This restriction will play an important role in understanding the quantitative discipline offered by the ATE.

What, then, are we looking for in the data to separate the relative importance of meeting and learning? For some intuition, note that because $\hat{\pi}$ and π are independent, the coefficient $\hat{\beta}$ is approximately

$$\hat{\beta} \quad \approx \quad \frac{\text{cov}[\log(\pi_i'), \log(\hat{\pi}_i)]}{\sigma^2_{\log(\hat{\pi})}} \quad \text{if } \hat{\pi}_i \geq \pi_i \text{ for all } i \in \mathbf{T}.$$

From this perspective, β therefore measures the covariance between a firm's *ex post* profit and its match's *ex ante* profit, normalized by the amount of exogenous variation fed into the experiment by way of the variance in mentor log profits.[15] Somewhat more intuitively, it tells us how match quality changes a firm's knowledge: If two treated firms meet with firm owners of different profitability but have similar *ex post* profit, the model infers that $\beta = 0$, or that no learning can occur.

Thus, the identification procedure boils down to the following: (1) check how much variation there is in treatment effects across individuals and how it varies with individual mentors, then (2) adjust the model-implied meeting rates so that the average of those individual-level treatment effects is consistent with the empirics.

In the interest of keeping the quantitative results self-contained, we hold off on estimating these parameters from our RCT until we discuss the remaining calibration of our full model below. We emphasize, however, that we could do so now and port those estimates directly into *any* model that satisfies Assumptions 1, 2, and 3.

We now turn back to our specific RCT and build a full general equilibrium model to study what we learn about the potential aggregate gains that could be achieved by scaling the program.

4 Full General Equilibrium Model of Diffusion

The RCT discussed above is a one-time, partial equilibrium shock to the source distribution. Our goal now is to understand what happens when the source distribution is permanently changed for everyone. That is, the impact of a policymaker changing the way in which people meet. This introduces two main complications relative to simply reading the results off the RCT. First, we need to consider the long-run implications. Second, we need to include general equilibrium effects that arise at scale. Therefore, we build out a general equilibrium model of knowledge diffusion

[15] The approximation is due to the finite sample. The exact equation requires orthogonalizing this relationship by projection on $\hat{\pi}$. That is,

$$\hat{\beta} = \frac{\text{cov}[A_{\log(\pi)} \log(\pi_i'), A_{\log(\pi)} \log(\hat{\pi}_i)]}{\sigma^2_{A_{\log(\pi)} \log(\hat{\pi})}} \quad \text{if } \hat{\pi}_i \geq \pi_i \text{ for all } i \in \mathbf{T},$$

where A is the relevant annihilator matrix.

that we can use as a laboratory to study these questions.

The backbone of the model are the components described in the previous section, and they will re-appear here. The goal of that section was to offer clarity on exactly how we move between parameters and empirical moments, stripping away any extraneous assumptions unrelated to identifying those diffusion parameters. Making quantitative statements, however, requires us to now make those additional assumptions. To do so, we close the model in a way to remain consistent with channels uncovered in the RCT. The economy is composed of individuals who decide between wage work or entrepreneurship, and then entrepreneurs have to search for and bargain with suppliers. This is nested inside a small open economy, but the labor market and diffusion of knowledge are local. Suppliers have access to infinite supply of inputs, consistent with their connection to the broader Nairobi and Kenyan economy. This broadly approximates the setting in which the RCT took place.

4.1 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, and her occupation. We refer to z as an agent's *knowledge*. The aggregate state of the economy is the distribution of knowledge, $M(z)$. Each agent dies with exogenous probability δ and a mass δ of new agents replace them each period. New agents draw their initial knowledge 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 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. We discuss these options in turn.

Profit, Knowledge, and Negotiation Motivated by our empirical results showing that the experiment in Brooks et al. (2018) caused lower costs and increased supplier switching, entrepreneurial knowledge affects profits through the ability to find suppliers with lower costs for the firm's inputs.

A firm purchases intermediate inputs x at a price p_x and pays a wage w to labor n. Production is Cobb-Douglas between intermediate inputs and labor, with weights α and η, respectively, yielding a profit function

$$\pi = x^{\alpha} n^{\eta} - p_x x - wn, \tag{4.1}$$

where the price of the firm's good is normalized to 1.

Markets for intermediate inputs in developing countries are typically decentralized, leading to search frictions, price dispersion, and price power. In order to purchase intermediate goods x, the entrepreneur must seek out a supplier. There is a continuum of suppliers indexed by their marginal cost m. If the entrepreneur purchases x units of intermediate goods at a price of p_x, the supplier's profit is

$$\pi^s(m,x) = (p_x - m)x.$$

Suppliers source their goods 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.

Entrepreneurs search for suppliers and then bargain with them. An entrepreneur with knowledge z exerting effort s matches with exactly one supplier whose marginal cost is

$$m = e^{-s} z^{\gamma}$$

where $\gamma < 0$. This specification implies that the elasticity of the marginal cost of the supplier that the entrepreneur meets is -1 and $\gamma < 0$ with respect to s and z, respectively. In order to render profits linear in z and not merely multiplicatively separable, let $\gamma = \frac{\alpha+\eta-1}{\alpha}$. This normalization is an ordinal transformation of the latent knowledge variable, and merely simplifies computation.

Thus, knowledge here is defined by its ability to lower the effort required to find low-cost suppliers.[16] The entrepreneur must find a new supplier each period (though the relevant knowledge to find the supplier depreciates more slowly).

Once they meet, the two parties engage in Nash bargaining over the price p_x that the firm will pay with bargaining weight ν. In equilibrium, they take into account the optimal input choice by the firm at that price for the profit function in equation (4.1), given reservation values of zero for both parties. This implies an equilibrium price

$$p_x^*(m) = \text{argmax}_{p_x} \, (\pi)^{\nu} (\pi^s)^{1-\nu} \tag{4.2}$$

that splits the surplus between the two parties. Thus, entrepreneurial knowledge enters the firm's profit function indirectly by reducing its intermediate input price $p_x^*(m)$.

Knowledge Transmission Knowledge is transmitted between agents. The learning technology that translates meetings into future knowledge is one we have assumed throughout. If an agent with knowledge z today meets another agent with knowledge $\hat{z}$, her individual stock of knowledge evolves as

$$z' = e^{c+\varepsilon} z^{\rho} \max\left\{1, \frac{\hat{z}}{z}\right\}^{\beta} \tag{4.3}$$

where c is a constant and ε is an exogenous shock drawn from a distribution with cdf F.

[16] The constant returns assumption for suppliers eliminates any strategic interactions between supplier-firm bargaining games. It guarantees that there are no issues of capacity considerations. See, e.g., Baqaee et al. (2024) on the macroeconomic consequences of supplier churn with firm-to-firm trade linkages.

We next define the source distribution $\widehat{M}$. While the literature typically assumes that the source distribution is identical to the firm knowledge distribution (denoted M^f and defined formally below), this need not be the case if agents intentionally seek out more useful imitation opportunities. In this case, the source distribution is better (that is, first order stochastic dominates) M^f. Alternatively, perhaps low z agents and high z agents are socially or physically distant from one another so that those high value imitation opportunities are much more rare than uniform. We nest these possibilities into a parameter θ, and assume as source distribution

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

The parameter $\theta \in (-\infty, 1)$ governs how different $\widehat{M}$ is from the existing knowledge distribution of firms, M^f. As the parameter θ increases, higher draws become uniformly more likely in the sense of first-order stochastic dominance. At $\theta = 0$, draws are uniform from M^f, but as $\theta \uparrow 1$, it concentrates all probability on the most knowledgeable entrepreneurs in the economy. Conversely, if θ is negative, selection trends toward less knowledgeable entrepreneurs. This simple transformation captures the cases discussed above. This parameter θ is the meeting parameter highlighted in the previous section, alongside the persistence parameter ρ and the diffusion parameter β, since it governs the extensive margin of diffusion. Notice also, per the discussion in Section 3, that we are required to assume a functional form for the source distribution. While we remain agnostic about the specific interpretation of the meeting parameter θ, Appendix B provides a variety of micro-foundations for this kind of specification.

Selection Into Entrepreneurship and Recursive Formulation With the entrepreneurial profit function defined, the flow utility of operating a firm is determined by

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

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. On the other hand, the flow utility for being a worker is

$$u^w = \omega \log(w)$$

where w is the equilibrium wage. The optimal occupational choice decision is then simply

$$u(z,M) = \max\{u^f(z,M), u^w(z,M)\}$$

and we denote decision rule $\phi(z,M) = 1$ as firm operation and $\phi(z,M) = 0$ as wage work. Taken together, the value of entering the period with knowledge z and aggregate state M is

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

where M' is next period's aggregate state.

Closing the Model and Equilibrium The occupational decision rule ϕ implies an equilibrium firm knowledge distribution

$$M^f(z,M) = \frac{\int_0^{\hat{z}} \phi(z,M)dM(z)}{\int_0^{\infty} \phi(z,M)dM(z)}$$

and therefore the underlying source distribution is

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

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 knowledge distribution $M(z)$, such that the value functions solve the agent's problem above with the associated decision rules, and the aggregate state evolves according to

$$\begin{aligned} M'(z') &:= \Lambda(M(z')) = \delta G(z') \\ &+ (1-\delta)\int_0^{\infty}\int_0^{\infty} F\left(\log(z') - \rho\log(z) - \beta\log(\max\{1,\hat{z}/z\}) - c\right)\widehat{M}(d\hat{z};M)M(dz). \end{aligned} \quad (4.4)$$

$\Lambda(M)$ is the law of motion for the aggregate state, and characterizes a stationary equilibrium in which $M^*(z) = \Lambda(M^*(z))$. It consists of the knowledge of new entrants, $\delta G(z)$, and the evolution of knowledge for the $(1-\delta)$ proportion of surviving agents.

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

Proposition 2. *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 2 turn out to be useful for calibration below.[17] The occupational choice margin creates an 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 D).

4.2 Model Calibration

We now turn to calibrating the model. This involves two steps. In the first step we estimate the diffusion parameters that govern learning (β,ρ) and meeting (θ). Per the results in Section 3 and Proposition 2, these link directly with RCT moments and are independent of the remaining model structure. We can therefore fix them during the second step, which is a more standard calibration. The full set of moments and parameter values are reported in Table 2.

4.2.1 Linking RCT Moments and Diffusion Parameters

Following the procedure laid out in Proposition 1, we first estimate

$$\log(\pi_i') = \tilde{c} + \rho \log(\pi_i) + \beta \log\left(\max\left\{1, \frac{\hat{\pi}_i}{\pi_i}\right\}\right) + \varepsilon_i \quad \text{for } i \in \mathbf{T}.$$

[17] That $\pi(z,w) = A(w)z$ follows because matching between suppliers and firms is deterministic. Uncertainty about 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.

using just baseline and wave 1 (1 quarter later) data. In our specific RCT, the fact that $\hat{\pi}_i > \pi_i$ for all i simplifies this to

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

The results of this first step are in Column (1) of Table 1, and we find that $\beta = 0.538$ and $\rho = 0.595$. The relatively low ρ implies 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. The result $\hat{\beta} = 0.538$ implies there is a strong, but far from perfect, internalization of match knowledge.

After checking that our (β, ρ) estimates imply the existence of a θ that can match the average treatment effect, we estimate θ. Letting $\mathscr{P}$ be the set of mentee-mentor pairs $(\pi, \hat{\pi})$ in the data and $m_\pi^f(\hat{\pi})$ be the empirical proportion of observed baseline profit levels for our representative sample of the economy. We can compute the sample counterpart to the population ATE:

$$\text{ATE}^{\text{model}}(\theta; \hat{\beta}, \hat{\rho}) = \frac{\frac{1}{|\mathbf{T}|}\sum_{(\pi,\hat{\pi})\in\mathscr{P}} \pi^{\hat{\rho}} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\hat{\beta}}}{\frac{1}{|\mathbf{C}|}\sum_{\pi\in\mathbf{C}}\sum_{\hat{\pi}\in\mathbf{C}} \pi^{\hat{\rho}} \max\left\{1, \frac{\hat{\pi}}{\pi}\right\}^{\hat{\beta}} \frac{1}{1-\theta}(m_\pi^f(\hat{\pi}))^{\frac{\theta}{1-\theta}}}. \tag{4.5}$$

Notice here that this makes use of the previous step (estimating $\hat{\beta}$ and $\hat{\rho}$), the observable mentor-mentee pairs $(\pi, \hat{\pi})$, and the fact that our baseline survey allows us to construct the profit distribution of firms in operation, m_π^f. As such, only θ remains unknown. We then determine the value of θ conditional on these other observable moments that equates $ATE^{\text{model}}(\theta; \hat{\beta}, \hat{\rho})$ with the empirical average treatment effect, given in Column (2) of Table 1. This adjusts θ until the equilibrium assignment mechanism delivers a distribution that rationalizes the observed difference between the treatment and control groups. This approach yields a value of $\theta = -0.417$. For a sense of magnitude, this implies that the expected match knowledge is 95 percent of the knowledge of the average operating firm. The average match is thus slightly worse than the expected draw from the operating firm distribution.

To re-emphasize the critical features here, deriving (4.5) relies on two specific assumptions on data and model. First, it exploits randomization between control and treatment. This allows us to eliminate unobserved shocks ε from both numerator and denominator. Second, it relies on Proposition 2, which allows us to translate unobservable knowledge into observable profit. Combined with estimate (β,ρ) it leaves θ as the only unknown on the right-hand side of (4.5).

4.2.2 Calibration of Remaining Parameters

We next calibrate the remaining parameters, holding fixed our previously estimated diffusion parameters. 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 learning function 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. 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 knowledge evolution are the parameters α and η, the knowledge growth term c, and the standard deviations of the exogenous shocks σ_0 and σ. The final parameter is the bargaining weight ν.

These 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.[18] The standard deviation of new entrant knowledge 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 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 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.

α and η are then set to match two moments. The first moment is the wage bill relative to intermediate spending. The average firm has a relative wage bill of 0.13. Our Cobb-Douglas assumption implies that $wn/p_x x = \eta/\alpha$, so we can set $\eta = 0.13\alpha$. Next, a consequence of Proposition 2 is that

$$\sigma_{\log(p_x)} = \left(\frac{1-\alpha-\eta}{\alpha}\right)\sigma_{\log(\pi)}. \tag{4.6}$$

The left-hand side is the standard deviation of log unit intermediate costs across firms. We find that $\sigma_{\log(p_x)} = 1.61$ in our baseline data after removing industry fixed effects. Since $\sigma_{\log(\pi)} = 0.99$ is also matched in the calibration, rearranging (4.6) after imposing these empirical estimates and $\eta = 0.13\alpha$ implies

$$\alpha = \frac{\sigma_{\log(\pi)}}{\sigma_{\log(p_x)} + 1.13\sigma_{\log(\pi)}} = \frac{0.99}{1.61 + 1.13 \times 0.99} = 0.36$$

[18] 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.

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

4.3 Dynamic Treatment Patterns

Before turning to quantitative results, we digress slightly to discuss the time series of the average treatment effect. Our calibration does not use the full time series of the average treatment effect, only the baseline data and the first survey wave after treatment. In this section, we examine how the model performs on subsequent waves to determine whether it actually captures the dynamics of entrepreneurial knowledge.

To do so, we run an identical experiment in the model and trace the time series of the ATE over 5 quarters. These results are in Figure 2a. 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 learning parameters 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.

One interesting result here is that higher β hastens fade-out. As we discuss below in the quantitative results, higher β is also critical for generating general equilibrium gains from diffusion. 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 shortly.

5 Quantitative Results

This section provides the general equilibrium implications of our model, disciplined by the RCT results. We start in Section 5.1 with our main quantitative experiment. Using the model developed above, we ask how the economy changes when there is a permanent change affecting agents' ability to meet others with high-knowledge. We show specifically how the quantitative magnitudes are informed by the RCT moments. In Section 5.2 we take on a more micro-development question. We ask how close our quantitative magnitudes are when we use only the average treatment effect from the RCT to discipline the model. Our interest here is in asking when, if ever, it makes sense to extrapolate the partial equilibrium average effect to the general equilibrium at scale effect.

Finally, we note that these results are all in pursuit of understanding changes to meeting technologies. But as these results will highlight, how people *learn* plays a central role in understanding them. We therefore take a step back in Section 5.3 and use the model to ask whether interventions that affect learning are more or less effective than those that affect meeting. While we do not these as closely connected to the RCT we ran, we use this to motivate future work in this area.

5.1 Scaling the RCT's Intervention

Quantitative Experiment To investigate the impact of improved availability of imitation opportunities in general equilibrium, we focus on a particular case and imagine increasing the meeting parameter θ so that it moves 25 percent closer to its limit of $\theta = 1$, and decompose the resulting change into the contributions of different general equilibrium effects. This increases baseline $\theta = -0.417$ to $\theta^{new} = -0.063$.

We think of this is as an aggregate policy change motivated by the RCT discussed earlier. The RCT finds promising results from a one-time change in meeting. We view this exercise as a policymaker being motivated by these results to institutionalize the ability of all agents to more easily meet with high-knowledge matches.[19] We study the long run, general equilibrium change in income that results by comparing the two stationary equilibria.

Results 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 knowledge distribution by making it easier to learn from high-knowledge agents. The second is an amplification effect through prices.

We decompose the relative importance of these two channels in columns 2 and 3. Column 2 isolates the direct effect. Average knowledge rises by 7 percent and the labor supply declines by 10 percent as the distribution shifts mass across the (fixed) cut-off ability level that defines occupational choice (see Proposition 2). Average income rises by 7 percent.

Column 3 allows the equilibrium wage to adjust. The wage increases by 13 percent as the knowledge required to find lower-cost suppliers increases the marginal product of labor. This competitive pressure causes the lowest knowledge 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 knowledge, 61 percent is from the direct effect and 39 percent is from

[19] The particular 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. For some sense of magnitude, this increase in θ increases expected match knowledge by 42 percent in the new steady state. In the Appendix we provide the results when the aggregate policy exactly replicates the estimated RCT gains. The same results hold with larger magnitudes.

the amplification through prices. A similar magnitude holds for income. Both forces play a quantitatively relevant role in generating the aggregate gains.[20]

5.1.1 Complementarity Between Meeting and Learning

We next turn to understanding how the model generates this 11 percent increase in income, and what factors play a major role. At the heart of our results is understanding the complementarity between meeting and learning. To see this, we start by measuring the quantitative impact of our policy shock (increasing baseline θ to $\theta^{new} = -0.063$) under various combinations of baseline meeting (θ) and learning (β, ρ) parameters. Throughout, we hold the remaining calibration fixed. These results are in Figure 3.

We start by focusing on the set of possible outcomes in the model. This set is shaded in gray in Figure 3. Depending on the parameters chosen, average income rises by anywhere from 0 to more than 4,000 percent. Complementarity between meeting and learning drives this result. All else equal, the gains are larger when it is more difficult at baseline to meet high-knowledge agents. That is, the larger the difference between baseline θ and $\theta^{new} = -0.063$, the larger the change in the meeting technology. But Figure 3a shows that even if θ is low, a wide range of aggregate outcomes remains possible. This is a function of the complementarity. If no one can learn ($\beta \approx 0$), there is no aggregate benefit regardless of θ. But if it easy to learn ($\beta \approx 1$), the aggregate gains can be extreme: over a 40-fold increase in income. The opposite also holds. As meeting becomes easier ($\theta \approx \theta^{new}$), the policy offers less benefit regardless of how easy it is to learn.

[20] Interestingly, this same change in θ allows the planner to increase consumption by only 6 percent in the efficient allocation, compared to 11 percent in the decentralized *laissez faire* equilibrium here. This is because the change in θ here decreases the magnitude of the diffusion externality. In the planner's allocation, that channel is absent because she already internalizes it. Thus, the benefits from better meeting are smaller. See Appendix D for more details.

5.2 ATE-Constrained Comparative Statics

Next we demonstrate the importance of separately identifying the intensity of diffusion β from the availability of imitation draws θ by showing that the average treatment effect from the micro level experiment taken by itself is consistent with a very wide range of aggregate effects in general equilibrium. Hence, it is very important to show *how* a given treatment effect arises. More formally, we re-run the identical experiment as in Section 5.1, where we increase baseline θ to be 25 percent closer to $\theta = 1$, but we do so under different combinations $\gamma = (\beta, \rho)$ with the restriction that the baseline θ satisfy

$$\text{ATE}(\gamma, \theta^*(\gamma)) = \text{ATE}^{\text{data}}.$$

That is, in the absence of our procedure, a given average treatment effect is consistent with a large set of parameters. For example, a high observed average treatment effect could arise either because diffusion is very strong (high β) or good imitation opportunities are rare in the control group (low θ). We show that the aggregate implications for the change in θ that we implement vary widely. The results are in Figure 4.

Figure 4a illustrates the set of parameter values consistent with our baseline average treatment effect. As β gets larger, the model-implied average treatment effect would rise. To keep ATE fixed, the model assigns a larger value of θ, since higher θ means the control group receives better imitation draws, which reduces the average treatment effect. Note also that when β is too low, the model can no longer rationalize the observed treatment effect.[21]

Figure 4b then plots the aggregate gains that are consistent with our observed ATE. The largest equilibrium gains (40 percent) are available at high β and low ρ in the bottom right of Figure 4b. The rationale here is that higher β and ρ both lower the implied θ, but β plays a larger role in

[21] That is, the ATE falls outside of the range discussed in Proposition 1.

aggregate outcomes than ρ. Thus, setting $\rho = 0$ and $\beta = 1$ creates the largest θ shock without sacrificing too much of the amplification from learning.

This large range of results demonstrates that the average treatment effect is not sufficient to be informative about the general equilibrium consequences of improved imitation opportunities. The results could be large or small for a given ATE. Our procedure for separately identifying these parameters is therefore crucial for determining where in this range the results lie.

5.2.1 When does covariance offer additional information about at-scale outcomes?

The previously results are derived conditional on average treatment effect we observed in our RCT. How general are these results? That is, if we instead observed a much smaller or larger average treatment effect, should policymakers come to the same conclusion that measuring β is critical to understanding general equilibrium gains? We show here that the answer to this question is "no." Only when the ATE is sufficiently large does our additional moment offer useful value. When the ATE is small, the aggregate gains will be small regardless of the other parameters.

To see this, we replicate our previous exercise and compute the band of possible aggregate outcomes for a given average treatment effect. But instead of using our observed ATE^{data}, we counterfactually vary the ATE between 0 and 100 percent and compute the bad of aggregate outcomes consistent with each. Those results are in Figure 5.

While a 1 percent ATE admits a fairly narrow band of aggregate gains between 0.01 and 0.05 percent, a 100 percent ATE admits gains between 10 and 285 percent. Thus, while a small ATE accurately predicts a small general equilibrium impact at scale, a large ATE offers a much wider set of possible at-scale outcomes. The intuition follows almost directly from the same complementarity highlighted above. With an ATE near zero, the model requires that either learning is extremely difficult ($\beta \approx 0$) or everyone can easily meet high-knowledge matches ($\theta \approx \theta^{new}$). Either of these implies no scope for aggregate gains. Thus, extrapolating from the ATE is warranted. On the other hand, a high ATE demands some combination of easier learning (β high) and worse baseline meetings (θ low). This opens scope for potentially large aggregate gains which is driven by the complementarity between exactly these two forces. Thus, extrapolating from the ATE is no longer warranted.

To see this slightly differently, Figure 5 shows that our 11 percent ATE is consistent with average treatment effects between 35 percent and 100 percent. In fact, our estimate of 11 percent is on the lower end of what could have been feasibly estimated given our treatment effect of 49.6 percent. In summary, measuring covariance turns out to be particularly important when interpreting evidence from RCTs that create large partial equilibrium benefits for the average participant. Moreover, it can be easily estimated in intervention-level data, making it a relatively straightforward moment to measure. On the other hand, covariance offers little additional information when experimental results show partial equilibrium average effect.

5.3 Comparative Statics on Diffusion Parameters

These results are all derived from changing θ, making it easier for agents to receive a high-knowledge imitation opportunity. But, at least in theory, one could imagine other RCTs directly affecting the ability to learn via changes in β. Such interventions would facilitate greater transfer of entrepreneurial knowledge from the more successful entrepreneur to the less successful one, and might be accomplished through more deliberate and structured interactions, as in a mentorship program designed to increase the flow of information rather than merely facilitate interactions.

Figure 6 provides comparative statics on these two parameters. Overwhelmingly, interventions that affect β have more significant impact on income relative to those that affect θ (so much so that the x-axis has to be different to see the patterns). Thus, our view here is that there is significant scope in economic environments with the potential for social learning to change income via learning interventions. We view this as a useful avenue for future work.

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 even conditional on the average effect. The results highlight the important complementarity between causal interventions and aggregate models (Buera et al., 2021a).

The results also open up additional questions for future work. We leave out features such as firms that are unwilling to share information due to competition, though we show in the Appendix that some versions of this idea are feasible under our framework. One could always write down a more complicated model with such a feature. That part is easy. But adding model features demands more empirical moments for estimation. This requires more subtly and targeted variation in designing interventions that speak to these features of the environment. 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.

Baqaee, David R., Ariel burstein, CEdric Duprez, and Emmanuel Farhi, "Supplier Churn and Growth: A Micro-to-Macro Analysis," 2024. Working Paper.

Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and A. Mushfiq Mobarak, "Can Network Theory-based Targeting Increase Technology Adoption?," *American Economic Review*, 2020. forthcoming.

Benhabib, Jess, Jesse Perla, and Christopher Tonetti, "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," *Econometrica*, 2020. forthcoming.

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, "Not All Management Training is Created Equal: Evidence from the Training Within Industry Program," 2019. Working Paper.

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.

_ , Joseph P. Kaboski, and Robert M. Townsend, "From Micro to Macro Development," 2021. Working Paper.

_ , _ , 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.

Caunedo, Julieta and Namrata Kala, "Mechanizing Agriculture," 2022. Working Paper.

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," *American Economic Journal: Macroeconomics*, 2022. forthcoming.

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*, 2020. forthcoming.

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

_ , "Recipes and Economic Growth: A Combinatorial March Down an Exponential Tail," *Journal of Political Economy*, 2022. forthcoming.

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 José Tessada, "Role Models or Individual Consulting: The Impact of Personalizing Micro-entrepreneurship Training," *American Economic Journal: Applied Economics*, 2018, *10* (4), 225–245.

Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E. Waugh, "The Welfare Effects of Encouraging Rural-Urban Migration," January 2018. NBER Working Paper 24193.

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*, 2007, pp. 3085–3113.

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," 2021, *111* (1), 73–128.

Romer, Paul M., "Endogenous Technological 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.

Van Patten, Diana, "International Diffusion of Technology: Accounting for Heterogeneous Learning Ability," 2020. Working Paper.

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.

Figures and Tables

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

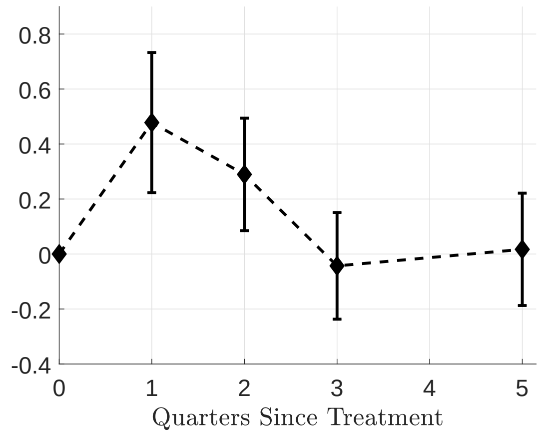

Figure notes: Figure plots average treatment effect as percentage above control mean (0.4 = 40%), along with the 95 percent confidence interval. Treatment takes place between quarters 0 and 1.

Table 1: Moments for Diffusion Parameter Estimation

	(1)	(2)
β	0.538	
	(0.273)**	
ρ	0.595	
	(0.273)**	
ATE^{data}		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, ***.

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 (3.2)	RCT results	0.538	0.538
ρ	Persistence of knowledge	0.595	Estimated parameter from regression (3.2)	RCT results	0.595	0.595
θ	Match technology "quality"	-0.417	Treatment effect over 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 knowledge distribution	1.00	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 knowledge shock distribution	0.73	Standard deviation of log profit in all firms	Baseline survey	0.99	0.99
c	Growth factor in knowledge evolution	-2.11	Mean profit over new entrant profit	Baseline survey	1.51	1.51
ω	Consumption utility weight	0.47	Fraction of employment in wage work	IPUMS	0.48	0.48
α	Knowledge elasticity in supplier search	0.36	Standard deviation of log inventory cost	Baseline survey	1.61	1.61
η	Knowledge 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.

Figure 2: Relationship between Treatment Persistence and Diffusion Intensity β

(a) Baseline Model vs Data

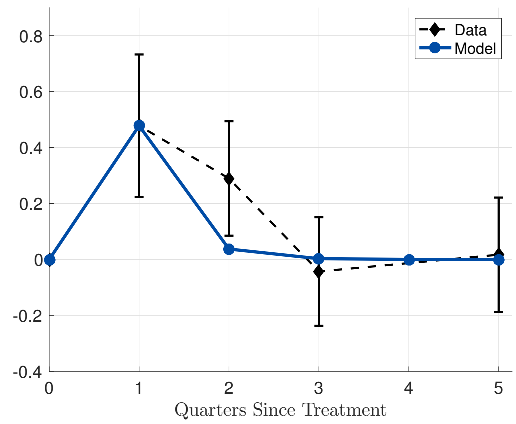

(b) At Estimated Persistence $\rho = 0.595$

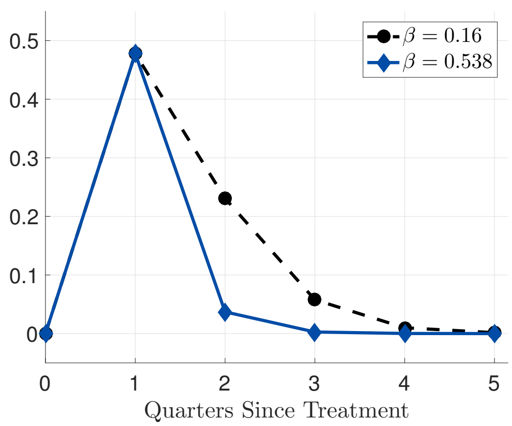

(c) At High Persistence $\rho = 0.99$

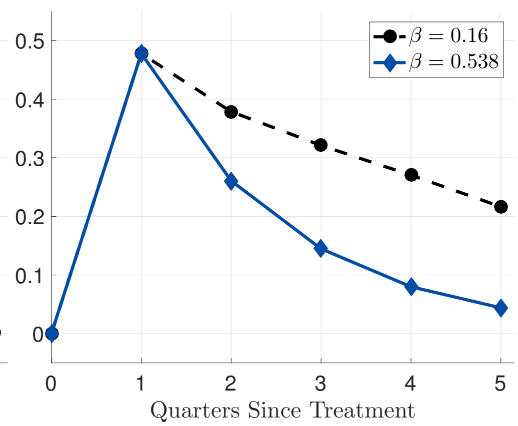

Table 3: Equilibrium Moments

	(1) Baseline	(2) Fixed Wage	(3) New Equilibrium
Income	1.00	1.07	1.11
Knowledge	1.00	1.07	1.12
Aggregate Wage-Labor Supply	1.00	0.90	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.

Figure 3: Possible Outcomes for Different Diffusion Parameters

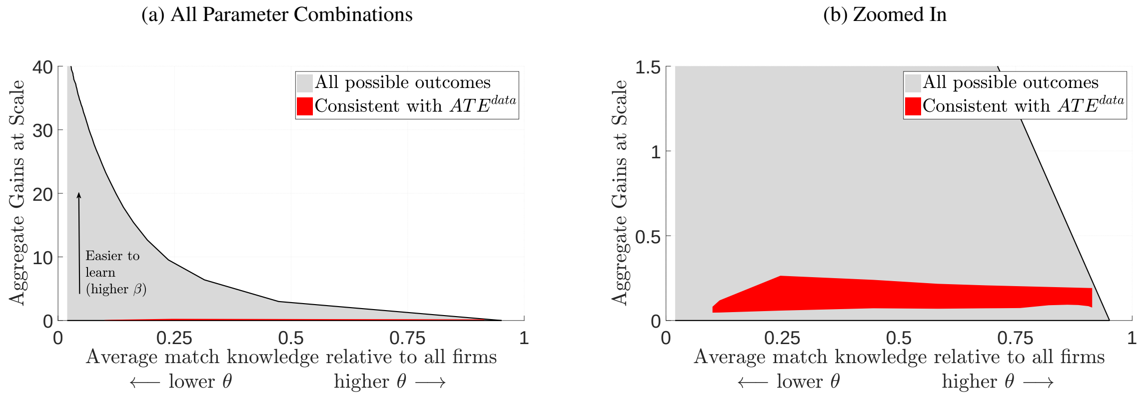

Figure notes: Aggregate income change after shock to $\theta^{new} = -0.063$. We consider the range $\theta \in [-250, -0.063]$ and $(\beta, \rho) \in [0, 0.99] \times [0, 0.99]$. Each economy holds the remaining calibration fixed at baseline parameter values. Figure 3b is a zoomed-in view of Figure 3a around the ATE-consistent outcomes. Multiply the vertical axis by 100 to get percentage gains (i.e., $5 \equiv 500\%$).

Figure 4: Aggregate Implications of Varying (β, ρ) at Baseline ATE

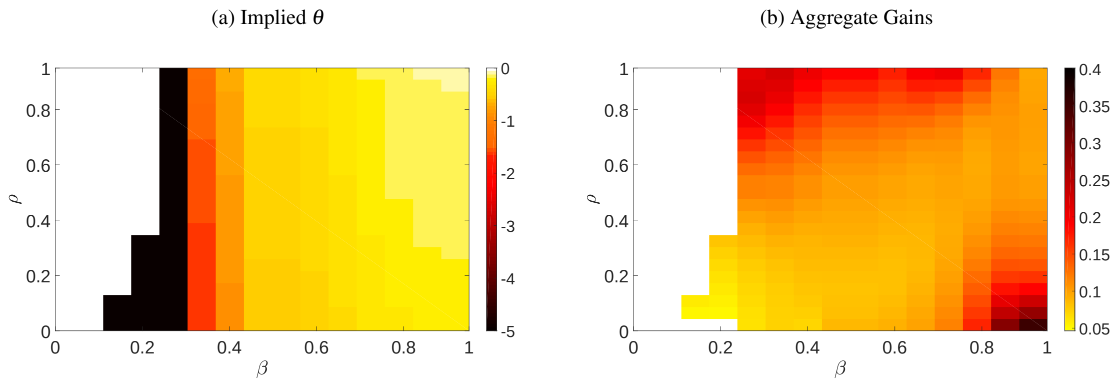

Figure notes: For each value of (β, ρ), (a) plots the implied θ that guarantees the model hits our baseline ATE, ATE^{data}. (b) then plots the implied percentage change in average income from better matching ($\times$100 for percent, i.e., 0.4 = 40%). Dark colors represent larger absolute values and white space means that it is not feasible to match our ATE at the given values of β and ρ.

Figure 5: Relationship between average treatment effect and equilibrium increase in income

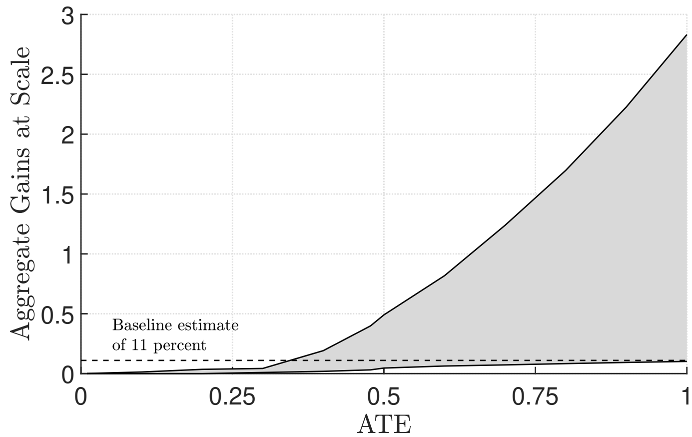

Figure notes: The shaded region shows all possible realizations of aggregate gains that can be achieved while holding the average treatment effect fixed by re-estimating the extensive margin parameter θ for each (β,ρ) combination. Multiply by 100 for percentage gains. The baseline ATE is $ATE^{\text{data}} = 0.496$.

Figure 6: Earnings Comparative Statics

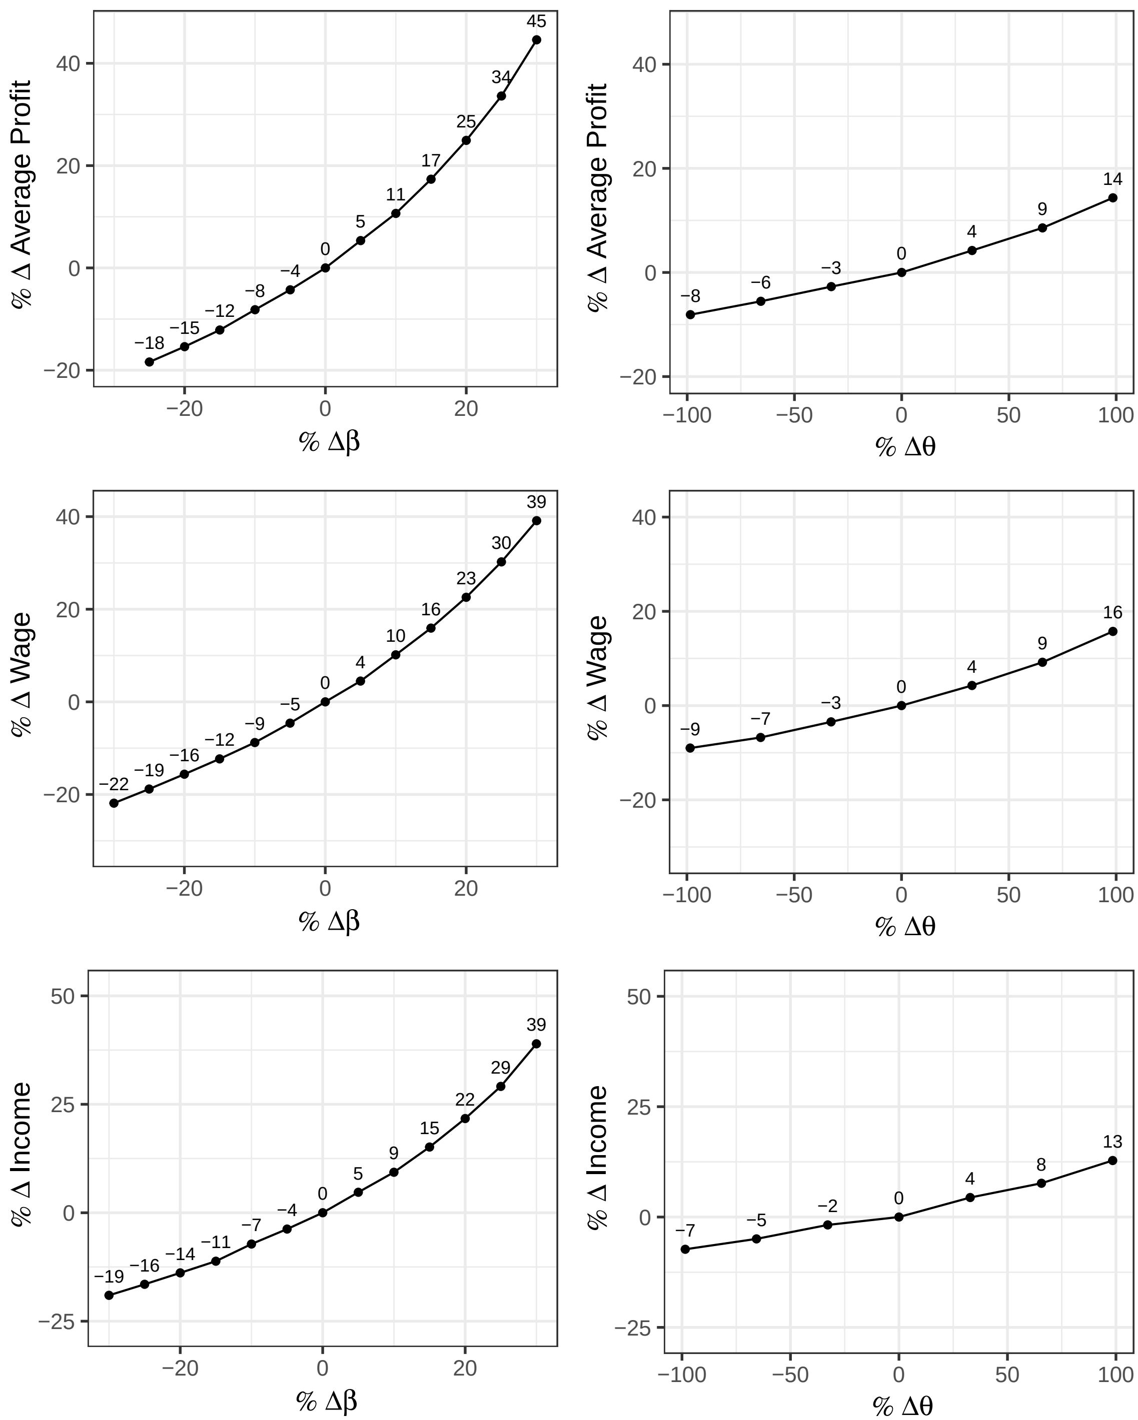

Table of Contents for Online Appendix

A Additional Details from the Main Text

A.1 RCT: Baseline Profit Distributions

Figure 7: Baseline Profit Distributions

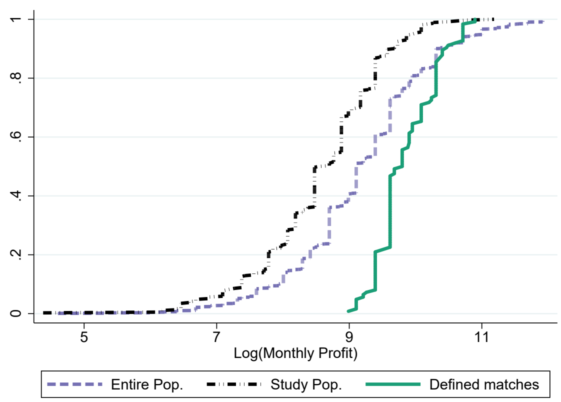

A.2 RCT: Balance Tests

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.[22] Table 5 reports the results. The only significant difference is in age, and the magnitude is small.

[22] 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, ***.

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 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.

Table 5: Balancing Test at Baseline

	Control Mean	T_L - Control	T_M - Control	T_H - Control
	(1)	(2)	(3)	(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.

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

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

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 flexi-

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

(a) Full sample

(b) IK Optimal Bandwidth

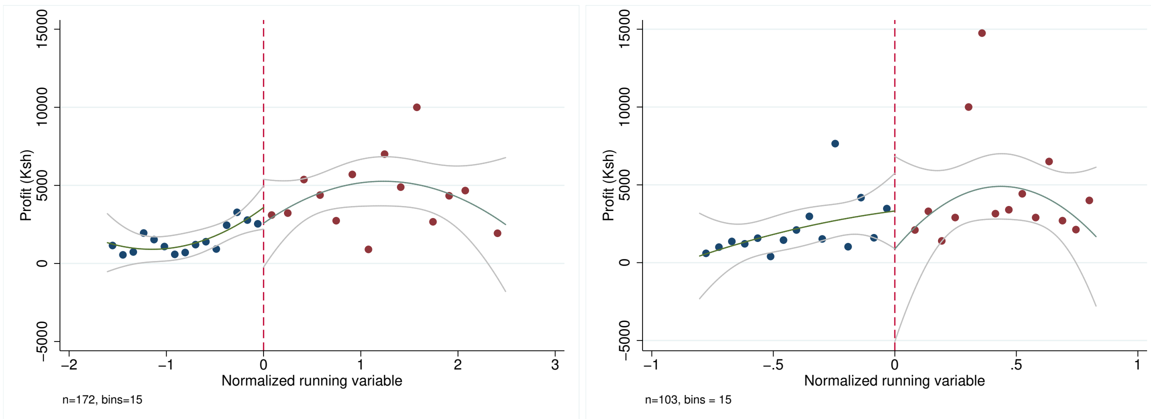

Figure notes: Figure 8 plots profit along with a fitted quadratic and its 95 percent confidence interval. Figure 8a uses the entire sample, while Figure 8b uses the Imbens and Kalyanaraman (2012) procedure to choose the optimal bandwidth. Both use 15 bins on either side of the cutoff.

ble 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 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 3.1), which is assumed here and in much of the existing literature.

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

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

Percent of IK optimal bandwidth	Scale		Practices	
	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 perc

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.[23] This implies that θ increases to $\theta = 0.512$.

Table 9 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 tec

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

[23] 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 $\widehat{M}$.

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.

B 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.

B.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:

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

B.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:

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

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:

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

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

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

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

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:

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

or following the notation more standard in the paper:

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

B.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:

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

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:

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

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:

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

B.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. (2020), who study learning among German co-workers.

C 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 E, we provide a quantitative evaluation of a particular extension related to mis-measurement.

C.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*

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

Assuming $f(\hat{z}/z) = \max\{1, (\hat{z}/z)\}^{\beta}$ gives us the original Assumption 1. Proposition 3 summarizes that we can instead estimate the function f using the same data-generating process as in the main text.

Proposition 3. *The data-generating process of Assumption 4 identifies (ρ, f) in equation (C.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.[24] Härdle et al. (2000) provides a detailed review.

C.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.

[24] 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.[25]

Proposition 4. *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.

C.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 E.

[25] 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:*

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

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.[26]

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

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

[26] 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.

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 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 5. *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 5. The key feature, however, is that this result allows a broad class of extremum estimators to be deployed to identify β.

C.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.

D 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 2, 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.[27]

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

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

[27] 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.

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

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

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

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

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

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 – (D.2) determines the resources that can be allocated to consumption, while (D.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 D.4) and the imitation opportunities that arise from it (via D.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 6. *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* $c_p = AN_s^{\frac{\eta}{1-\alpha}} Z^{\frac{1-\alpha-\eta}{1-\alpha}}$*, where A is a constant and the two aggregate* *inputs are*

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

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 G. ■

One can see the goals of the planner in Proposition 6. 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.[28] We derive these results in Appendix G.3, and proceed with quantitative results in the next section.

D.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

[28] 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 $s(\underline{z})$.

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 is normalized to one at our estimated value.

Figure 9: Importance of Separately Identifying Two Diffusion Effects

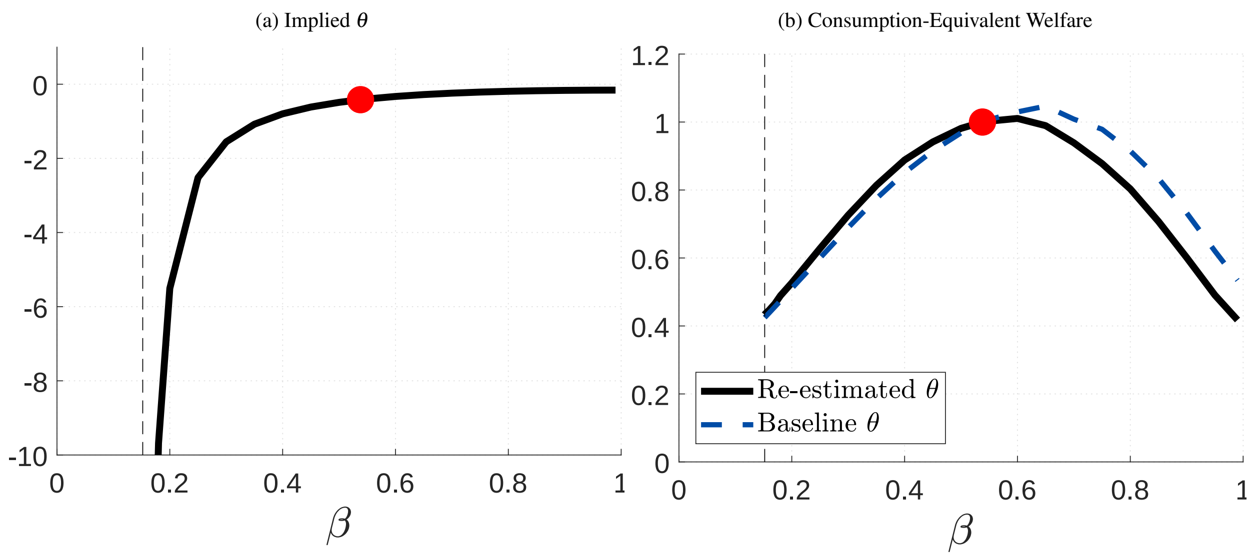

Figure notes: The left panel shows the implied θ required to estimate the same treatment effect observed in the data at the exogenously given β. The dashed vertical line the minimum value of β for which there exists a θ that can rationalize the observed average treatment effect. $\theta \to -\infty$ 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.

D.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.[29] 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 6, 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,

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

Equation (D.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.[30] Figure 10 plots the semi-elasticity ε_{c_p}, along with those of the two aggregate inputs defined in (D.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:

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

[30] 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

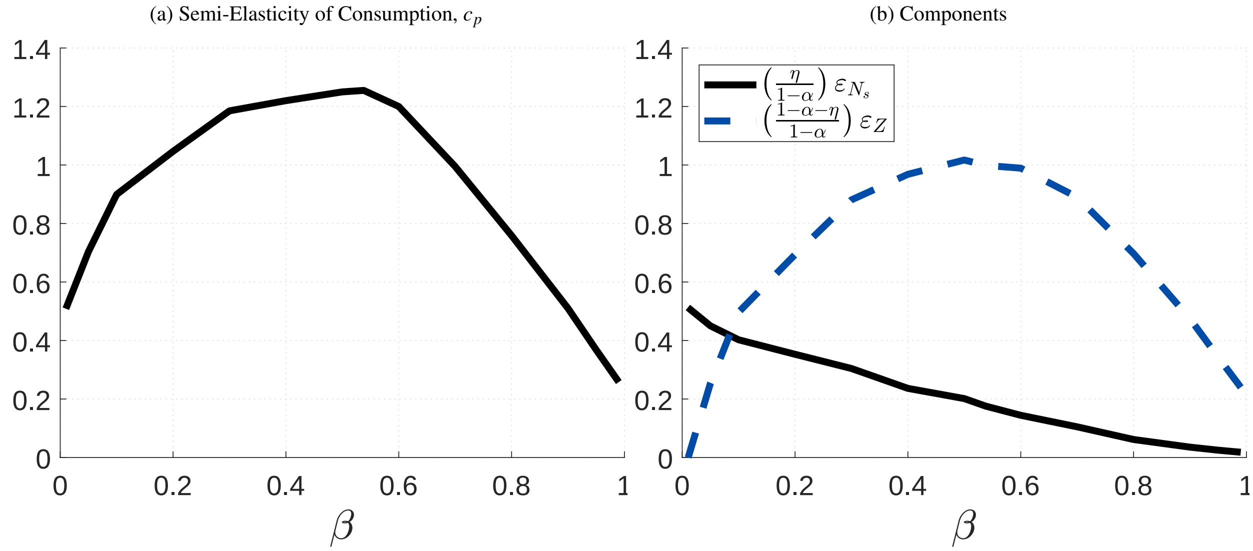

Figure notes: These semi-elasticities are evaluated at the baseline equilibrium cut-off $\underline{z}^*$ and equilibrium distribution $M^*(z;\underline{z}^*)$.

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 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 2 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

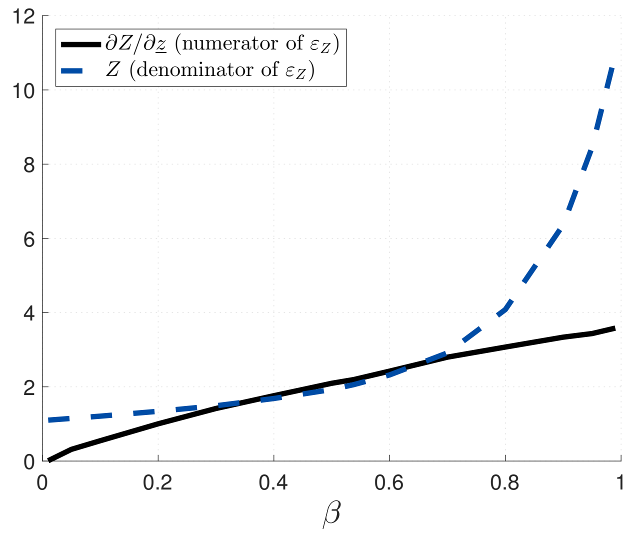

(b) Wage w and Cut-off Ability $\underline{z}$ in the Baseline Equilibrium

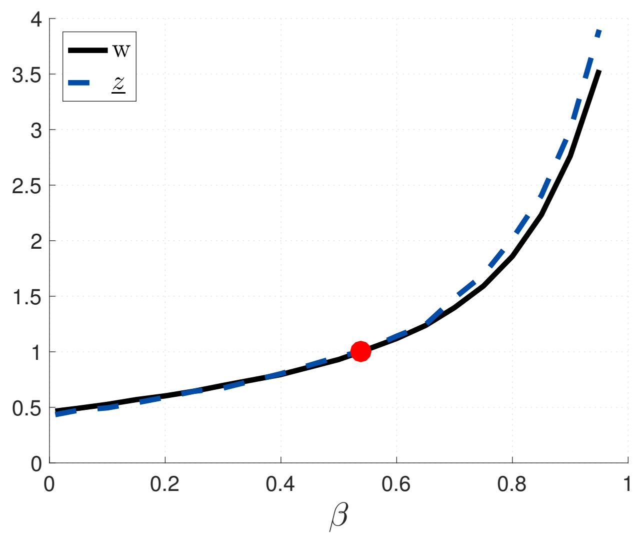

(c) Average Consumption

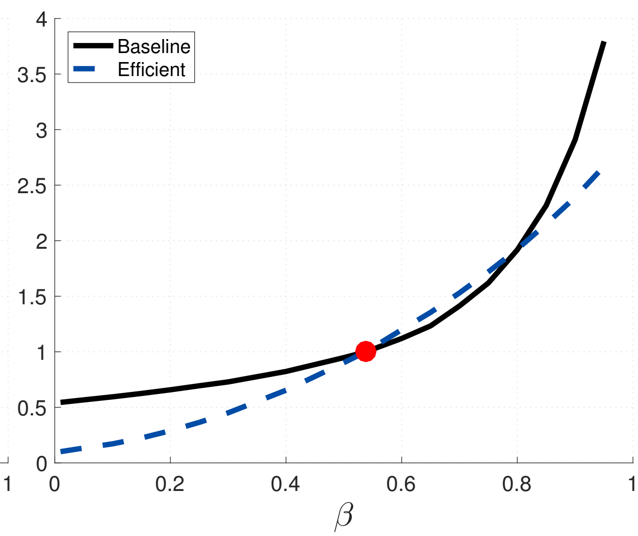

Figure notes: In Figures 11b and 11c, our estimated values are denoted by the circle in each graph and normalized to one.

E 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 C this can be extended to estimate the distribution of τ. We leave this aside for simplicity here. Define the function

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

where $\blacksquare := (c, \rho, \beta)$ is the parameter set. Our first stage regression is $\log(\pi'^*) = g(\pi^*, \hat{\pi}^*; \blacksquare)$, 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.

E.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.[31]

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 $\hat{\pi}$), we

[31] 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.

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

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

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 $\blacksquare := (\tilde{c}, \rho, \beta)$ to solve

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

where

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

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

E.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 8, along with the baseline for comparison. We do so in two scenarios: $\sigma_\tau = 0.30$ and $\sigma_\tau = 1$.

E.3 Quantitative Exercise

We now study the gains from the same exercise as in the text. The main results are in Table 9. 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 8: 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

Table 9: 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

Figure 12: Range of Aggregate Gains for each ATE

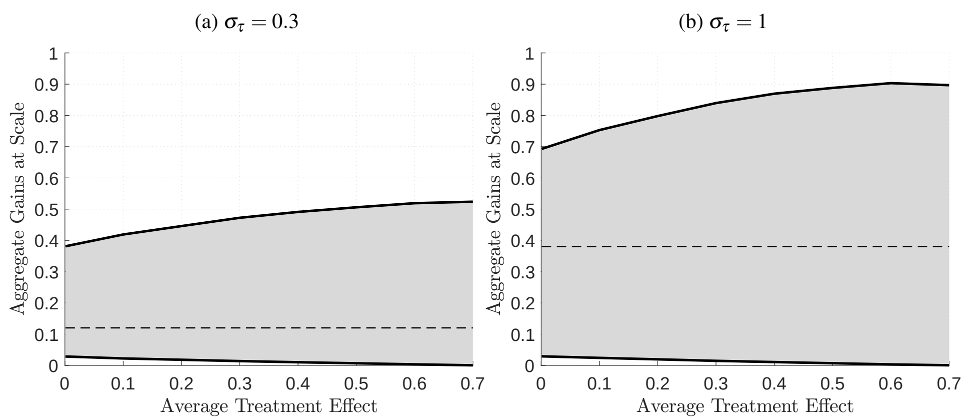

Figure notes: Shaded area is all possible aggregate gains for $(\beta, \rho) \in [0.20, 0.95] \times [0.20, 0.95]$, where θ is chosen to match the ATE on the horizontal axis. The baseline estimate is given by the dashed line.

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

F.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.[32] 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 $t = 2$.

F.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. $\lambda_4 > 0$).[33]

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

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

A number of additional results provide context for these firm performance effects. First, they find persistent changes in management practices.[34] 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 $\bar{n}_i^m$ as the average firm size of firm i's matches, they run

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

They find statistically significant increases in sales and profit. CS use (F.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.

F.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

[34] 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.

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 replaced by δ new firms, who draw initial productivity from $z \sim G(z)$.[35] 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:

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

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

[35] 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.

a match, it gets $\hat{z} = 0$.[36] We can therefore write 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

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

Equation (F.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

[36] 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.

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

F.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 C.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

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

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

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

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

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

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 (F.5) as the average treatment effect

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

where $T_i = 1$ if i is in the treatment group. Table 10 provides the regression estimates for (F.4) and (F.6).

Table 10: 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.

F.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 11.

Table 11: 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 match 1-1 with target moments. Parameters in Group 3 are jointly calibrated.

F.6 Treatment Effect Persistence

In the main text, we showed that the treatment effect persistence is 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)

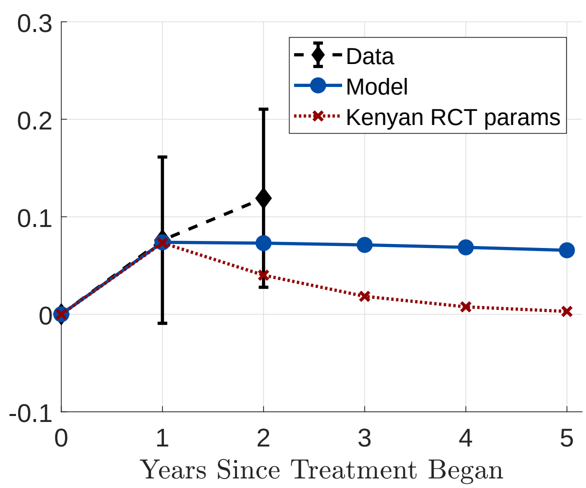

Figure notes: Dashed line is the data with 95 percent confidence interval, for which data is available at $t = 0, 1, 2$. Solid line is the estimated model $(\beta, \rho, \theta) = (0.276, 0.955, 0.199)$. For comparison, we include the results at our Kenyan RCT values of β and ρ, re-estimating θ to hit 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.[37] 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.

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

F.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 12. 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 12: 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 matchi

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.[38] 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.

[38]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 1 for more details on this constraint.

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

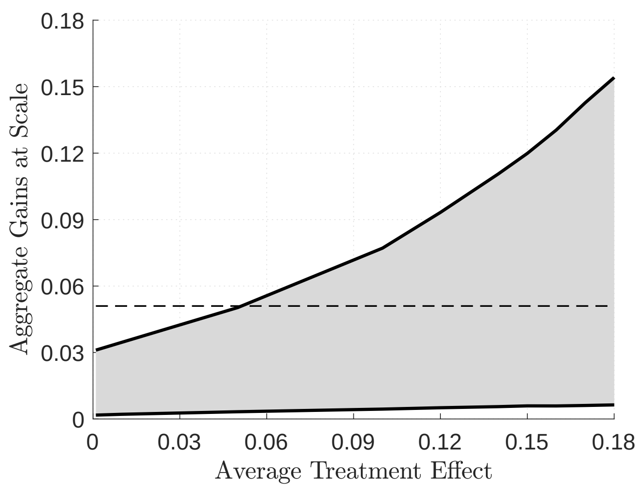

Figure notes: Shaded area is all possible aggregate gains for $(\beta,\rho) \in [0.20, 0.95] \times [0.20, 0.95]$, where θ is chosen to match the ATE on the horizontal axis. The baseline estimate at $(\beta,\rho) = (0.276, 0.955)$ is given by the dashed line.

G Proofs

G.1 Proof of Proposition 1 (estimation strategy)

Recall that the bounds used in Proposition 1 are given by

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

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

Proof of Step 1: Unique (β, ρ)

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

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

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, (G.1) is a simple panel regression with coefficients β and ρ. That $\hat{\beta}^{OLS}$ and $\hat{\rho}^{OLS}$ are equal to their structural counterparts follows directly from the fact that this regression has full rank (Assumption 4 (c)) and that unobservable shocks ε are orthogonal to π and $\hat{\pi}$ (Assumption 4 (d)). ■

Proof of Step 3: There exists at most one θ **for each** (β, ρ) **that matches the ATE**

Proof. In the model, we can write $\pi' = g(z, \hat{z}, \varepsilon) = Ae^{c+\varepsilon} z^{\rho} \max\{1, \hat{z}/z\}^{\beta}$ 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."[39]

[39] 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.

Inserting the correct joint distributions implies

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

Applying the proportionality assumption in Assumption 2 and the orthogonaility condition of Assumption 1 gives us

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

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.

■

G.2 Proof of Proposition 2 (equilibrium characterization)

We begin by detailing the underlying arithmetic of the model, then use those results to prove Proposition 2 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

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

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

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

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

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

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

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

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

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

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 2 With these results, Proposition 2 follows quickly.

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

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

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 (G.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. ■

G.3 Proof of Proposition 6 (from Social Planner's Problem)

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

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

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

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

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 (G.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.

■

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*, 2020. forthcoming.

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.